# euROBIN Strategic Research Agenda

*Sponsored by:*

**euROBIN Network of Excellence**

THE EUROPEAN EXCELLENCE NETWORK
ON AI-POWERED ROBOTICS

This project has received funding from the European Union's Horizon Europe Framework Programme under grant agreement No 101070596.

info@eurobin-project.eu @euROBIN_project @euROBIN-project www.eurobin-project.eu

# TABLE OF CONTENTS



## ANNEXES

### 1. Will AI alone solve robotics?



### 2. Control, Planning and Reasoning in the era of generative AI



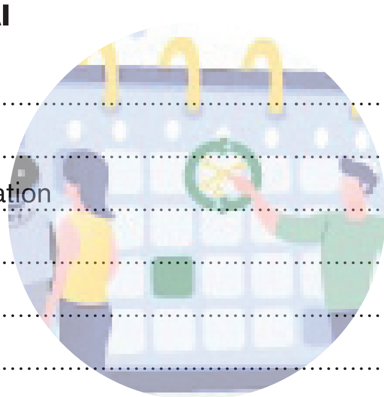

### 3. Human-Robot Interaction: Successes, Hurdles and Remaining Challenges





www.eurobin-project.eu
info@eurobin-project.eu
@euROBIN_project
@euROBIN-project


euROBIN website
(Featuring EuroCore and SRA)

euROBIN Week 2024
Video of the 1st euROBIN Coopetition @Nancy


This project has received funding from the European Union's Horizon Europe research and innovation programme under grant agreement No. 101070596

# EXECUTIVE SUMMARY

*"AI has been advancing at an incredible pace. It started with perception AI — understanding images, words, and sounds. Then generative AI — creating text, images and sound. Now, we're entering the era of physical AI, AI that can proceed, reason, plan and act."*

Jensen Huang, Nvidia founder and CEO, CES 2025, January 5, 2025

Robotics has experienced a transformative evolution over the last two decades, progressively expanding its reach far beyond its origins in industrial automation. Robots today are significantly more autonomous, more intelligent, lighter, more robust, and less costly than all previous generations. These achievements have been driven by advances in modelling, dynamics, control, planning, learning, and reasoning, sensors and actuators, computing power, miniaturization of electronics and batteries, 3D printing and additive manufacturing, materials and especially soft materials. But robotics is actually today not only a "decathlon of engineering and computer sciences", but it benefits from highly interdisciplinary research, involving biomechanics, neuroscience, medicine, psychology, and industrial design, in addition to all the engineering disciplines. The convergence of these developments with rapidly evolving artificial intelligence techniques has created the foundation for a new generation of cognitive, adaptive, multi-purpose machines - commonly referred to as **Physical AI, Embodied Intelligence,** or AI-Powered Robotics. At its core, intelligent robotics fundamentally manifest AI in physical form.

The strategic importance and potential of intelligent robots extends beyond technical innovation; it is central to address some of humanity's most pressing challenges, as highlighted by the United Nations' Sustainable Development Goals (SDGs) and the major strategic pillars of the Horizon Europe program. A major challenge which is addressed by robotics, is the demographic change, implying a decreasing and ageing population. The shrinking working-age population puts pressure on labour markets and the economic competitiveness, healthcare system, and welfare states. In Europe there are 22.2 million people working in human health and social work. The new wave of AI-based robotics can provide, in the next decade, massive relieve in this respect. Moreover, manufacturing represents the first economic driver in Europe, accounting for over 80 % of industrial value added and more than 31 millions of employees, followed by water supply, sewerage, waste management and remediation activities[1,2]. Robotics can greatly contribute to give impetus to these sectors that require intelligent physical action and robots to substantially support humans in all these activities. Robotics can successfully address the challenges of effective, sustainable and resilient industry, efficient and environmentally-friendly agriculture and food chain supply, effective asset predictive maintenance and management, as well as the growing challenges on security, safety and healthy working and life conditions. Realizing this vision demands sustained investment in fundamental research, cross-disciplinary collaboration, and strong engagement between academia, industry, and policy makers. A comprehensive strategy must align Europe's expertise in AI, robotics, and engineering with real-world application areas such as healthcare, advanced manufacturing, agriculture, space exploration, and environmental sustainability.

[1] *https://www.statista.com/statistics/1195197/employment-by-sector-in-europe*

[2] *https://ec.europa.eu/eurostat/cache/htmlpub/key_figures_on_european_business_2021/industry.html*

The robotics landscape is at an inflection point, driven by hardware advancements and by AI breakthroughs such as deep learning, reinforcement learning, and foundation models. While, until recently, even the largest robotics companies were of moderate size in the overall economic landscape and most robotics research was performed by academic and research institutions, we are now witnessing a fundamental transformation in the field. Major U.S. technology companies - including Tesla, Google, Amazon, and Nvidia - are making massive investments in robotics, while new players such as Figure and Physical Intelligence have emerged with billion-dollar funding rounds. The declared ambition of many of these companies is to achieve general-purpose physical intelligence, thus aiming at creating a humanoid robot capable of performing most, if not all, physical tasks that a human can accomplish. While expectations and the market hype surrounding this vision are undoubtedly exaggerated in the short term, the level of investment and talent being directed towards embodied AI clearly indicates that this next stage in robotics is here to stay. Even if a true breakthrough will take a decade or more, it is clear that robotics is entering a new era of accelerated development. Mirroring this trend, China has designated humanoid robotics as one of its core technological priorities, alongside electric mobility, renewable energy, and low-altitude transportation. **The question is how do we respond to this paradigm shift?**

Robotics is a domain where European research and industry are at the forefront of international developments, see the next section regarding the state of the art. The European Commission through its various funding schemes together with national programs in European countries have led to a world-leading research environment, a strong tradition of industrial and service robotics, and a well-established ecosystem of suppliers. These are very strong assets as the new robotics race unfolds. A cohesive and forward-looking strategy—uniting academia, stakeholders and industry— is essential to ensure leadership in robotics.

A **balanced robotics AI strategy** must invest not only in algorithmic development but in the full ecosystem where AI will operate - **including physical platforms, sensing technologies, processing, and actuation systems.** Some of the world's leading AI players already recognize that intelligence cannot exist in isolation from the physical world, coining the term **Embodied AI** to highlight the deep integration of AI with robotics. While many focus on training ever-larger models on gigantic datasets, true AI leadership requires advancing embodied intelligence - where perception, reasoning, and action coalesce in real-world environments. Achieving this vision demands:

- **Scalable and cost-effective robotic hardware** for sensing, computing, and actuation, with a focus on resource conservation, life cycle extension, carbon footprint, energy efficiency, circularity and recyclability.

- **Building up and expanding capabilities in transferable AI methods including foundation models,** especially in the forms relevant for robotics, such as Visual-Language Models (VLM) or **Visual-Language-Action Models (VLA).**

The road to general-purpose AI in the physical world will be shaped by both near-term and long-term technical challenges. Beyond advances in language models and computer vision, robotics requires breakthroughs in modeling, control, planning, learning, reasoning, and human-machine interaction and integration - not just at the cognitive level but also at the physical level. Critical hardware technologies—including actuators, sensors, processing, batteries, and novel materials - must evolve in parallel with AI algorithms. A major concern in both AI and robotics research is that excessive specialization in niche models may limit broader cross-disciplinary innovation.

The **euROBIN** network therefore identified at its start in 2022 that achieving cognition-enabled transferable embodied AI is the most fundamental scientific and technological challenge that is currently hindering the breakthrough of AI-powered robotics, and hampering its wider deployment and commercialization. The main scientific aim of **euROBIN** is therefore to advance robotics by providing robots with general, flexible, and pragmatic methods for the transfer of skills and knowledge among robots, tasks and environments. Robots are general-purpose machines in principle; endowed with an advanced ability to transfer solutions across different tasks, domains, and robots, they will also become general-purpose machines in practice. The recent international developments mentioned above confirm that the forecast and focus were clear-sighted, even though foundation models – a novel popular approach to achieve transferability – were not known at project launch. The topic of transferability in the robotics context is so demanding that it will continue to be a main focus for research and development over the next few years.

To build truly **general-purpose robots,** systems must be capable of interacting with unfamiliar, unpredictable environments, learning new tasks from humans, other robots, and their surroundings, and operating autonomously in unstructured settings. Many of these environments, from disaster zones to deep-sea exploration, lack predefined models or sufficient training data, necessitating data collection, active perception, and lifelong learning strategies. These challenges will define the next frontier of AI-powered robotics.

To ensure that the next generation of robotics serves society, a **human-centered perspective** must guide development, integrating economic growth, sustainability, climate resilience, and ethical considerations in line with the SDGs. Technological progress must be accompanied by research into the **social, ethical, and legal** dimensions of intelligent robotics. Building **trust in AI-powered robots** requires understanding how attitudes toward robotic systems evolve across cultures and over time. Embodied AI must be designed not merely for efficiency but for **fairness, transparency, and social equity.**

To preserve a vantage position in robotics, leading in both research output and industry creation, and consolidate this leadership, it will be crucial to retain control and strengthen the entire value chain, from design to manufacturing and deployment and innovation. Past mistakes in critical technology sectors, such as semiconductors, have led to dependencies that must not be repeated in robotics. The window for decisive action is narrow.

A strategic approach to face the unfolding robotics race is to leverage our European diversity, the abundance of knowledge and the capabilities distributed across the continent. **euROBIN**, for example, addresses this challenge by building a common software and data infrastructure as well as a cooperative way of advancing and benchmarking robotics capabilities. This allows us to advance jointly at a faster pace, making use of the potential of a vibrant community. The goal is to bring many of the larger research labs to an end–to-end robotics competence level in major application fields and to better disseminate the in-depth competence and results of specialized labs across Europe. New start-ups and technology transfer to larger companies are very likely to emerge if we sustain and extend this activity over the upcoming years. Maintaining and raising leadership in the new era of robotics will require **massive further investment, a coherent strategy, and a coordinated effort across policy, research, industry, and the financial sector. The time to act is now.**

## ROBOTICS ACHIEVEMENTS IN THE LAST DECADES

Research on robotic manipulators, historically the first large family of robots, saw significant achievements in the last 20-30 years. Robots can manipulate and grasp objects with different shapes, with varying positions and trajectories. Compliant control allows them to be guided by the forces applied on them rather than being programmed to reach predefined positions, thus adjusting to contact with objects or humans. Steady progress has been made in the design of new materials that led to several prototypes, including a few commercial products of soft grippers and soft robotic hands for manipulating non-rigid objects, and proofs of concept of soft actuators fluidic, pneumatic, or based on functional materials such as shape-memory alloys. In the medical field, researchers showed that it is feasible to harness neural or muscular signals to control prosthetic limbs. In addition, digital twins of large environments enable robot programming at a more abstract level and facilitate the realization of robotics tasks of greater complexity, for example in production lines.

Transfer of these and other technologies to the commercial domain has resulted in industrial robots that can tackle increasingly complex tasks. A key advancement has been the introduction of collaborative robots (co-bots) that have expanded automation into less structured industry environments. Several robotic hands and grippers are on the market, for manipulating and grasping different categories of rigid objects. Surgical robotics has grown since the early 2000s into one of the main markets for service robotics, and many European start-ups are now leveraging miniaturization, machine learning and soft components to lower prices and expand use-cases. Wearable robots augment human bodies. Robotic arms operate on spacecrafts and autonomous rovers, demonstrating robustness in an extremely challenging environment.

Allowing robots to operate autonomously in novel situations and to approximate the dexterity and agility of living organisms have been key challenges for robotics since at least the 1960s. For several decades, robotics researchers have been experimenting with neural networks and machine learning as a potential solution to those challenges, and there is now sizable literature on how to leverage these techniques to tackle robotics problems that had previously proven hard to solve. Today, techniques to teach robots still rely on the two principal styles of machine learning that have been employed in robotics since the 1990s. On one side there is a family of algorithms that allow robots to learn from expert data, typically provided by a human demonstrator who demonstrates the target action while their movement is captured by visual or motion sensors. This approach has proved applicable in tasks ranging from grasping to manipulation of complex objects. The other type of learning algorithms enables robotic systems to learn through trial and error without a prior formalization of what constitutes the correct control policy. Best exemplified by reinforcement learning (RL) this method typically relies extensively on computer simulations of the robots and its environment to create enough learning cycles and learn a robust enough policy before testing it on the actual robot. Use of RL in robotics was hindered, for a long time, by the exploration phase, which, if not properly bound, can become too computationally and time intensive and its inability to easily scale to high dimensions. Recent advances leverage the increasing effectiveness of deep-learning and visually realistic physics-based simulation.

The design and control of autonomous vehicles, from cars and drones to walking and swimming robots, to enable navigation in different environments has improved thanks to sensors such as compact IMUs (Inertial Measurement Units), compact cameras, compact LiDAR, and to the use of machine learning. SLAM (Simultaneous Localization and Mapping) is now a mature technology (mainly originated in Europe) that allows mobile robots to explore an unknown environment while building a map of it. Legged robots can now walk on regular or moderately irregular terrains. Humanoids have become more agile and can bend down to pick up objects on the floor and walk while carrying objects in conjunction with humans or other robots. Aerial robots have robust flight capabilities and various degrees of autonomy, including autonomous flight at high-speed relying on cameras when in controlled light conditions. Drones with manipulation capabilities have been developed and tested, bridging the gap between the two traditionally separate areas of robotics. Underwater robotics offers a wealth of solutions, from small robot fishes to larger machines that can explore the depth of oceans. A few prototypes of autonomous cranes and other large machines for construction have been deployed.

Advancements in autonomous navigation have resulted in the commercial success of robotic vacuum cleaners (now the largest share of marketed commercial robots), lawn mowers and pool cleaners. Partially or totally autonomous drones are now being deployed for surveillance, monitoring, defense, inspection, and maintenance. The ongoing real-world tests of driverless cars are also a result of decades of research on autonomous navigation (again, first pioneered in Europe). Wheeled robots are used in logistics and warehouse management, and mobile robots of various types are applied to search and rescue and disaster management. Finally, wheeled and tracked rovers have been sent to explore bodies of the solar system.

Since the birth of modern robotics, European robotics research and technological development have been (and currently are) major drives on a global level, underpinned by a strong manufacturing and service robotics industry. Programs of the European Commission such as Horizon Europe, H2020, EIC, and ERC have played a pivotal role in fostering scientific

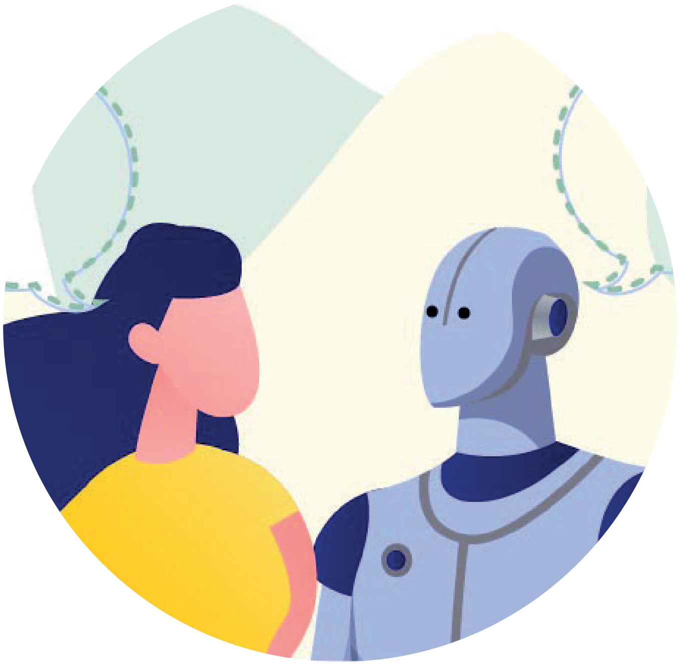

excellence and collaboration and technology transfer across Europe and strengthening international partnerships. European researchers contributed pioneering achievements, such as the first remotely operated robot deployed in nuclear service, Mascot (1959), still operational today in the Joint European Torus, or the first long-distance autonomous car by Dickmanns in 1986.

Today, Europe is the worldwide leader in robotics academic research. Over the last 5 years, EU countries published 54819 papers in Robotics, surpassing China (39912) and the United States (36160). Together with associated countries (UK, CH, TU, IL), the count is 72600  (source: Scopus).

Europe is also an important robotics market. With over 90,000 industrial robots installed in 2023, EU countries are second only to China (276,000) and in front of Americas (55,400) . 778,000 industrial robots are in operation in Europe, which corresponds to 18% of the worldwide robot stock.(source: IFR World Robotics Report on industrial robotics). In the service robotics market, Europe is leading. 405 companies (44%) are located in Europe. Asia (268 companies) holds a share of 29% and 233 companies (25%) are from the Americas (almost exclusively North America) (source: IFR World Robotics Report on Service Robotics).

European-funded robotics research has not only produced fundamental results but has also significantly impacted the development of our industry. One successful example of technology transfer from EU-funded programs to industrial development was the birth of the collaborative robot industry (including Universal Robot, Franka Emika, and dedicated ABB, KUKA, and COMAU new product lines). This was made possible by the development of a new generation of light-weight, sensitive manipulators with advanced control features and the scientific assessment of safety risks for humans. These results emerged to a large extent from EU-funded projects such as Phriends and solidified in the relevant ISO norms on collaborative robotics. This has resulted in a sector that today is worth 10% of the overall industrial robotics market and is growing. Similarly impactful research breakthroughs mainly originated in Europe are probabilistic robotics with its impact in Simultaneous Localization and Mapping (SLAM), which was a major enabler of mobile robotics. Many European breakthroughs have opened entirely new research and application areas, including Robot Programming by Demonstration, Soft Robotics, Bioinspired Systems, Neuromorphic control, and Neural interfaces for prosthetics and rehabilitation.

Robotics research advanced and accumulated over the last 20-30 years the ingredients necessary for a large leap step. preparing to enter the era of AI-based general-purpose robots. It is now time, also for Europe, to take the next steps and harvest the results of this long-lasting research endeavour.

# EUROBIN'S PROPOSAL FOR EUROPE ROBOTICS STRATEGIC AGENDA

Equipped with the appropriate level of knowledge, robots will have a certain level of understanding of the world that surrounds them and be capable of interpreting its dynamics in real time and reason about it. Deep learning, large-language models, and other AI technologies have gone from one breakthrough to the other. However, action and sensing in the physical world pose different and greater challenges than analysing data in isolation. AI-powered robots will be endowed with faster control capabilities for real-time planning and reasoning, robust execution of actions capable of safely handling unexpected contacts and disturbances, and the ability to optimize energy consumption during motion and task execution. New bodies and hardware components enabling robots to act and perceive the world as humans, and to become more efficient, will open new avenues for the deployment of robots. In parallel, advances in materials, micro and nanotechnologies, biohybrid systems open new avenues for importantly improving interaction, interfaces and integration between the human body and the robotic system, thus enabling more effective communication, collaboration and cooperation. New manufacturing processes, such as additive manufacturing and 3D bioprinting, will provide a great boost to robot fabrication processes by addressing the huge challenge of deploying affordable technology and sustainability.

The key ingredients to achieve this are multi-faceted. We divide them here in five major categories ("Robot Learning and Reasoning", "Planning and control", "Human-robot interaction", "Human-robot integration": Bionics and Biohybrid Robotics", and "Robot bodyware"). This document provides a brief overview of the key challenges facing AI-powered robotics from both research and technology perspectives and examines the challenges coming from applications in multiple domains. The document has been edited by the **euROBIN** consortium, based also on an extensive set of interviews involving worldwide robotics leaders, to include a global perspective. The discussions additionally materialized in a set of articles co-authored by those experts, and which address the challenges of the fields above at a higher level of scientific detail, and which were submitted to high-rank robotics journals.

## ROBOT LEARNING AND REASONING

The unprecedented advances in foundation models, from Large Language Models (LLMs) to Vision-Language-Action models (VLAs), has raised expectations that general intelligence for robots will become a reality. The vision of such robots - capable of understanding high-level instructions, decomposing them into executable tasks, and adapting dynamically to different environments - suggests a profound shift in how disembodied AI and robotics interact. Robotics presents unique challenges that cannot be solved by solely resorting to data-driven approaches.

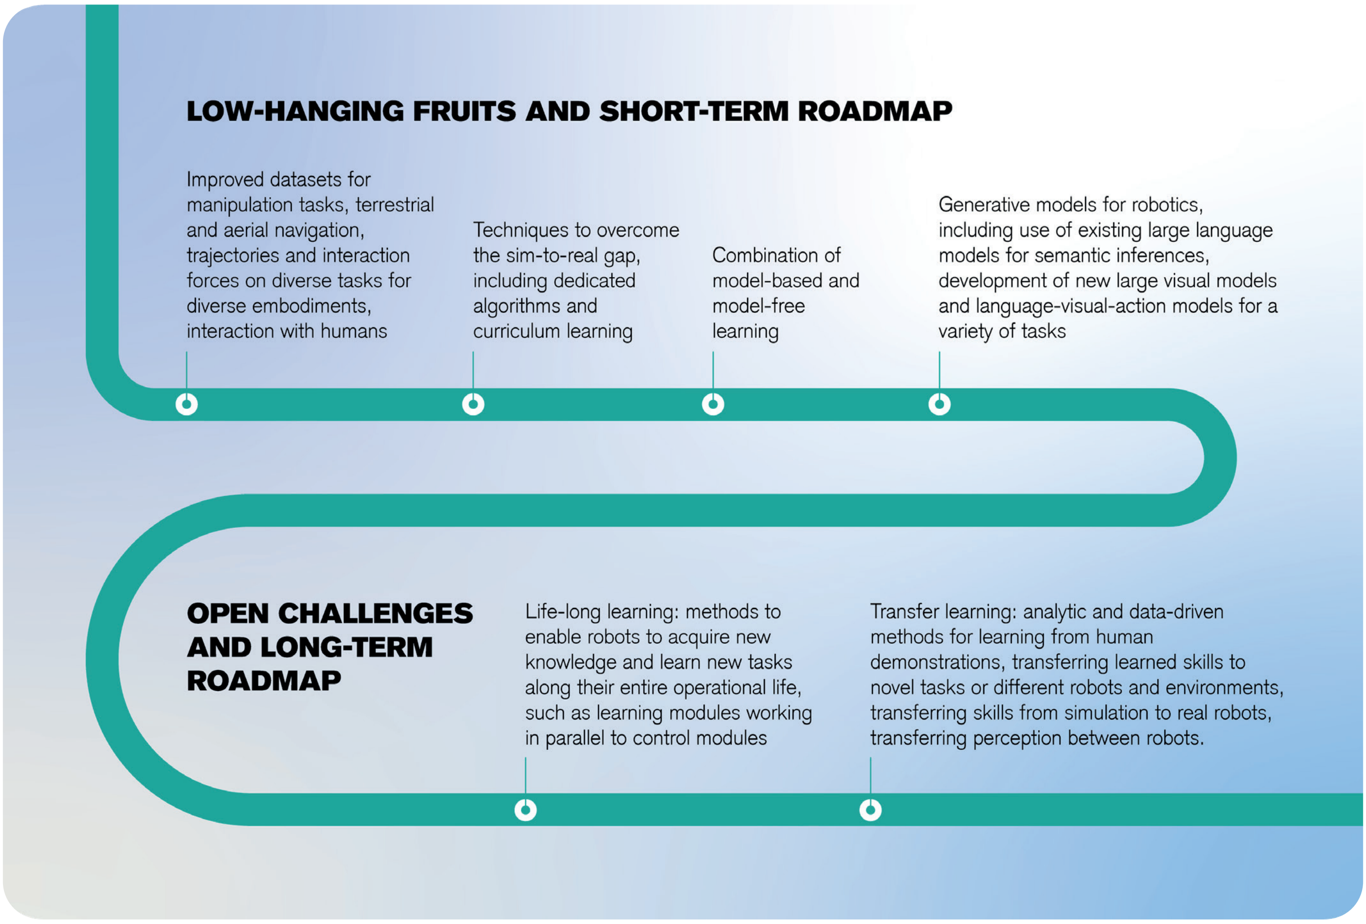


*Figure 1: Short-term and long-term challenges for further endorsement of AI in robotics, order by increasing level of complexity. Note that these challenges may not be overcome sequentially and proceed in parallel.*

For example, whereas text-based AI can be trained from data collected at scale from the internet, robot learning requires grounding high-level reasoning, understanding semantics and causality in **real-world physical interactions, safety constraints, and multimodal sensorimotor experience** - aspects that cannot be fully captured through passive data collection alone. Addressing these challenges demands a strategic combination of **model-based reasoning, structured learning, and data-driven adaptation** rather than relying solely on learning to discover solutions that are already well understood through physics and engineering principles.

One fundamental challenge is **bridging the reality gap** between simulation and real-world environments. While learning in simulated environments can accelerate development, transferring policies to the real world remains difficult due to imperfection of dynamic models, sensor noise, and difficulty modeling environmental variations. Even with domain adaptation techniques, purely data-driven models often fail to generalize beyond the specific conditions they were trained in. Research must focus on improving sim-to-real transfer through hybrid approaches that integrate data-driven learning with physics-based models to allow robots to refine and validate learned behaviors in real-world conditions. This will ensure that robotic control strategies remain robust across different environments and tasks. A second critical challenge is ensuring **safety and generalization** in robot learning. Unlike language models, incorrect predictions and failures in robotics can lead to physical damage, system malfunctions, or even harm to humans. Research must develop methods that allow robots to **predict and mitigate risks before executing**

**actions.** This includes human-in-the-loop safety mechanisms, explainable AI techniques and formal verification methods to create learning robot systems that are not only data-efficient but reliable. Combining model-based control with learning-based adaptation will be crucial to ensure that robots can generalize across tasks without compromising safety.

The issue of **data scarcity** also presents a major research challenge. Unlike text and image datasets, which can be collected at scale, robotic learning is constrained by the **high cost and slow acquisition of physical interaction data.** Future research should explore ways to maximize **knowledge transfer between robots, environments, and tasks,** reducing the need for extensive real-world data collection. This includes developing methods for transfer learning, lifelong learning, and meta-learning, allowing robots to continuously adapt based on past experiences and shared knowledge. Additionally, improved digital twin technology should be explored to enable more abstract, scalable programming of robots in simulated environments before deployment.

One of the most pressing open questions in robotics AI is **closing the signal-to-symbol gap** - the challenge of converting sensorimotor data into structured, abstract, symbolic representations that support reasoning, planning and decision-making. Current learning methods are effective in recognizing patterns in data, but they are limited in generating explainable knowledge for high-level decision-making. To address this, research must prioritize **neurosymbolic AI approaches that combine neural-based learning with structured reasoning** as a pathway to enabling robots to **interpret their actions, predict consequences, and adapt dynamically to new situations thanks to learned scalable, interpretable and generalizable models across environments, robots and tasks.**

Another critical research direction is **ensuring explainability and trust in AI-powered robotics.** Many real-world robotic applications require transparent, interpretable decision-making, particularly in fields such as healthcare, autonomous driving, and industrial automation. Unlike black-box deep learning models, which offer little insight into their internal logic, robot systems must be designed to provide justifications for their actions, detect when they are uncertain, and allow for human intervention when needed. Research should focus on developing AI frameworks that integrate knowledge, causal reasoning, and symbolic AI methods to enhance interpretability while maintaining adaptability. A further critical aspect for massively data-driven AI methods is their huge energy consumption, in particular in the training phase. The potential economic benefit of more resource-efficient AI methods is so large that it will trigger massive resource investment in this direction. First results, such as those of DeepSeek, attracted a lot of attention and fast evolution is expected in this area.

Figure 2: Handing a package from a drone to a humanoid robot or single-arm robot manipulator requires to reconcile drastically different perception, from different viewpoints and sensors, and distinct robot actions from unimanual to bimanual actions

Addressing these challenges requires a shift towards **hybrid AI architectures that blend learning, modeling, and reasoning,** ensuring that robots are not only capable of performing tasks but also of flexibly executing multiple and increasingly complex activities. They should understand and adapt to their environments in a meaningful, safe, and explainable way. By combining **model-based engineering principles with data-driven adaptation,** the next generation of intelligent robots will be able to generalize across tasks, transfer knowledge efficiently, and operate autonomously in real-world settings without sacrificing safety or reliability.

The future of AI in robotics will likely depend on a balanced integration of model-based approaches and learning-driven adaptation. Structured physics models, symbolic reasoning, and hybrid neuro-symbolic AI will play an essential role in enabling robots to reason, plan, and act safely and reliably in unstructured environments. Another key challenge is investigating bio-inspired models grounded in neuroscience to bridge the gap between artificial learning and execution for adaptive control.

Finally, some remarks on the very fast evolving field of foundation models for robotics: A quick succession of increasingly powerful and multi-modal GenAI models – **Large Language Models (LLMs), Vision Language Models (VLMs), Vision Language Action models (VLAs)** – have started to blend the lines between language, vision and action. By doing so, they have demonstrated impressive transfer of knowledge and skills between tasks at different levels of abstraction. LLMs can now provide robots with access to a wealth of knowledge about the world in linguistic form; a feat unthinkable only three years ago, even after decades of hand-crafting ontologies for this purpose. Large Reasoning Models (LRMs) extend LLMs with approaches such as Retrieval Augmented Generation and Test-Time to Compute to improve reasoning abilities, which is essential for robust robotic task planning. VLMs are able to accurately annotate (parts of) images with descriptions of astonishing detail and depth, whilst seamlessly recognizing and parsing text in the images with the same model. VLAs have demonstrated state-of-the-art skills on difficult tasks such as cloth folding. Even if many of these models are made open source, the developments are so rapid that benchmarking these models in the robotics community to understand their capabilities and limitations lags behind their releases. Nevertheless, it is clear that GenAI will be an essential component for achieving general-purpose robots.

An open research question for robotics in this context is how to integrate GenAI with typical robotics modules such as forward models, dynamics models, and memory-based world models. We know such modules to be essential for high-performance robotics, but they are not easily embedded in current GenAI models such as transformer networks. As current VLA models already have separate modules for vision, language and action to make learning tractable, we expect modularity will remain essential for more complex robotic architectures. A key question is thus how to develop **hybrid architectures in which the generality of GenAI approaches is combined with the safety, reliability, and efficiency of model-based approaches to reasoning, planning, control, and human-robot interaction.**

## PLANNING AND CONTROL

As robot bodies become more complex, new solutions are emerging to extend the methodical basis for modeling and controlling them, with tools from geometric mechanics and dynamics, differential geometry, and algebraic topology. At the same time, machine learning will play an increasing role in robotics, especially for systems for which physical models are lacking or are not accurate. The trade-off between model-based control and machine learning is one of the main challenges of future robot control. Developing **real-time deployable AI-based control** models that have an analytical interpretation is a major goal on this path.

The mathematical modeling of soft systems must be improved, to effectively **control soft robots** and their interactions with deformable objects. Both machine learning approaches and analytical models of manipulation of using geometric mechanics and dynamics, differential geometry, and algebraic topology, that can mathematically describe highly nonlinear dynamics will be developed in the next years; Similar methods can be used for models of complex interaction with fluids, such as air and water, which are crucial for many applications, from medicine to agriculture to the automation of the textile and food sectors.

The classical approach of fast, reactive, but local control and slow, global motion planning will be overcome as the boundary between the two methodological fields will vanish. Advances in algorithms, computing power as well as the use of machine learning make global planning real-time capable, while advances in differential and topological formulations of kinematics, dynamics and control open up the perspective of global task level and configuration space control.

Reasoning methodologies should evolve to incorporate motion planning and task-planning, interwinding control and planning over a long temporal period and moving towards a symbolically specified goal, a mission rather than merely a target position. Cobots can benefit from increased customization and more plug-and-play solutions based on combined reasoning and planning. They will include advanced interaction skills that will make it easier for end users to adapt them to different tasks.

**Multimodal** locomotion **control** is needed to allow legged or amphibious robots to change modes of locomotion, switching from walking to jumping, climbing, squeezing through narrow passages – including not only efficient control of each different gait but also, crucially, autonomous decision on when to switch gait. Legged systems such as quadrupeds and humanoids will largely benefit from a co-design of hardware and control, using massive parallel simulation in AI-based evolutionary processes. Advances in the understanding of resonances on elastic multi-body nonlinear systems will lead to robots with considerably increased versatility and energy efficiency.

Control of flying robots can be enhanced to combine stable flight with manipulation, including on-board real-time perception and planning. Significant improvements are needed on control of **multi-robot systems** where several robots of different types co-operate on tasks, including control of robotics swarms with several tens, or even hundreds, of individuals: here, progress will be required on creating shared representations of the environment (that different robots may observe from varying points of view and with different sensors) as well as on common operative systems and communication protocols.

Significant improvements are also needed on the real-time and computation performance of the control methodologies, to dramatically increase the performance and reduce the deployment efforts of the applications.

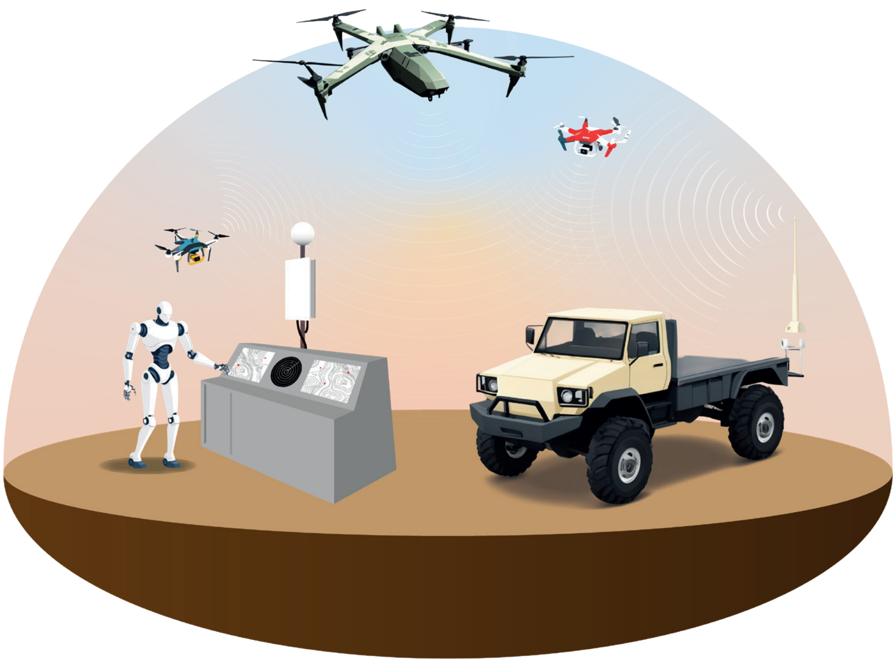

Figure 3. An open challenge is how to control multi-robot systems where several robots of different types co-operate on tasks and share representations of the environment. that they may observe from varying points of view and with different sensors.

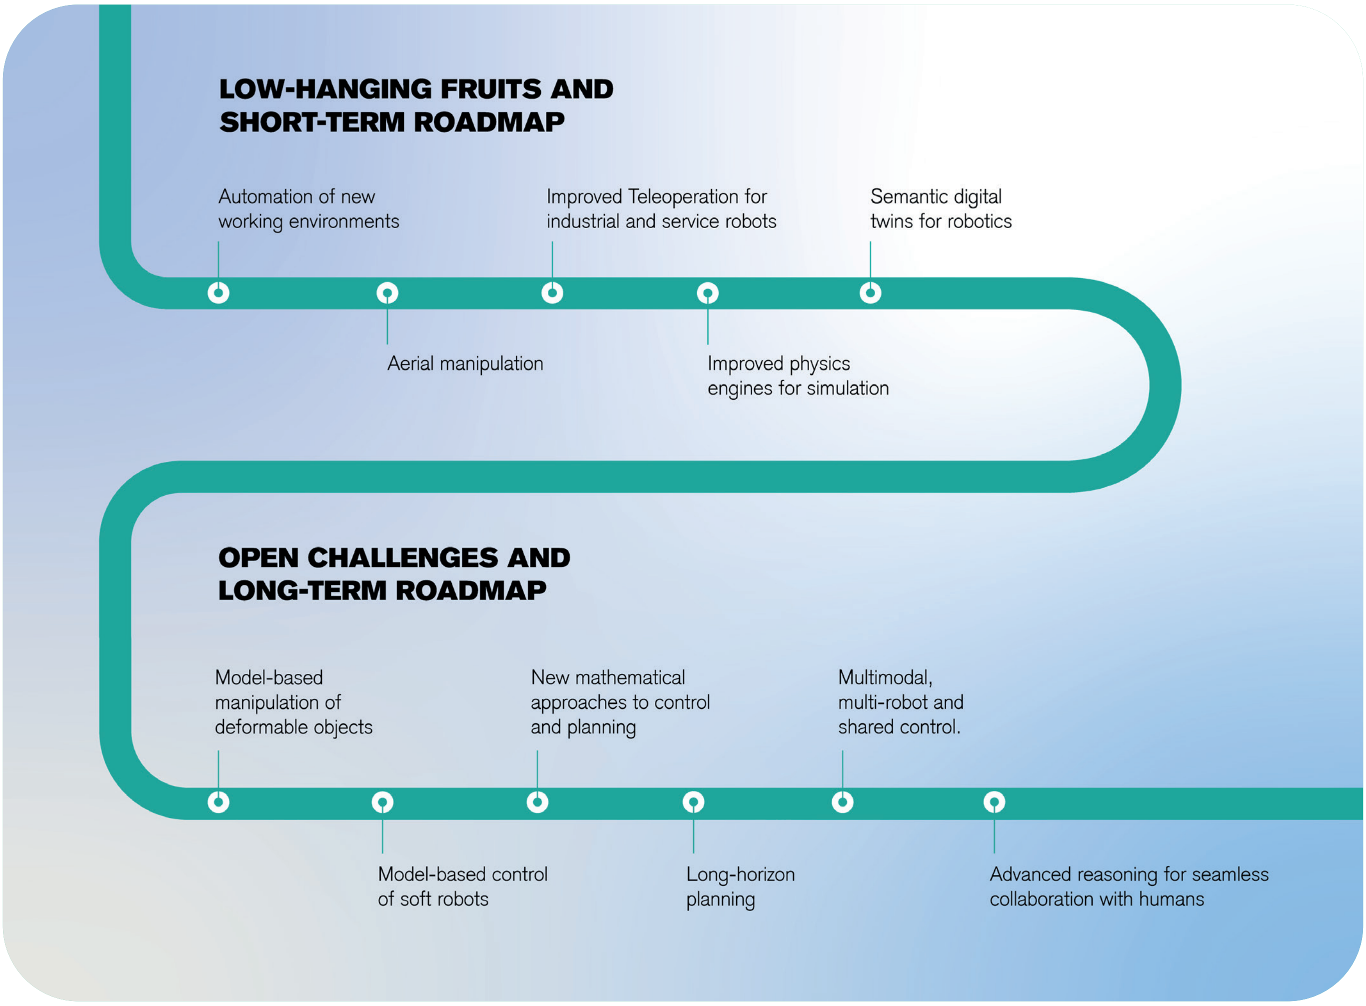


Figure 4. Summary of short-term and long-term research goals for control, planning and reasoning. The roadmap is not intended as a temporal sequence, but rather as a series of goals with increasing levels of complexity to be researched in parallel.

## HUMAN-ROBOT INTERACTION

The past decades have seen an increasing number of robots deployed in the vicinity of humans, from workers' companions in manufacturing settings, vacuum cleaners roaming in our living rooms, drones flying over our heads, to prostheses attached to our bodies. To increase trust and reduce risks, it is urgent and necessary that robots become **cognizant of their environment and socially aware.** They must be able to interpret, predict and reason about both human behavior and their own behavior.

Future AI-powered robots should explicitly account for human actions, preferences, mental states, and goals, factoring in privacy and related laws, enabling them:

- To determine when to act or communicate effectively
- To recognize when to assist, such as when a human is unwell, and stepping back upon recovery
- To adapt and give priority to the human, allow them the freedom to make their own decisions, and assist rather than impose their working rhythm.

Safe and intuitive human-robot interaction for non-expert users must be achieved, including two-way communication with robots based on either **physical interaction**–when humans and robots get into actual contact with each other, as in the case of rehabilitation and assistive robots, or **non-physical**, relying on verbal or non-verbal interactions, e.g. natural language, emotion recognition and control/perception skills. **Human-centric strategies to control** robots while interacting with the individual and the environment are also required, relying on dynamic adaption to the user's needs, multidimensional user monitoring, multisensory integration.

Gesture and speech recognition as well as AI-based body interfaces can be applied on a larger scale to give commands to robots. Improved interaction between humans and soft robots is going to be another medium-term research focus, including wearable devices capable of real-time state estimation of the body thanks to soft sensors.

Shared representations of the environment that different robots may observe from varying points of view and with different sensors are required. Integrating task planning with real-time feedback, robots can effectively co-construct actions with humans, ensuring mutual understanding and efficiency.

Ubiquitous perception - achieved by sensing objects, subjects, and gestures - represents a crucial challenge to achieving productive human-robot interaction. In addition to vision-based sensors, RF-based sensing technologies could be disruptive in the field. The ubiquitous perception involves a thorough analysis of ethical constraints and privacy risks.

Future challenges are:

- Striving to develop robots that do not overly constrain humans, by improving the intuitiveness, and ergonomics of human-robot interfaces, to facilitate their adoption
- Developing controllers capable of not only maintaining appropriate distance from humans, but also understanding human movements under different circumstances, even when they are confused, to ensure safe navigation in crowds and other heavily populated areas as well as safe collaboration in manufacturing and industrial settings.
- Conducting broader assessments to evaluate robot interactions with multiple users, ensuring controllers do not exhibit bias against minorities, with a particular focus on protecting vulnerable populations.
- Scaling up the use and deployment of collaborative robots, both outside and inside industrial settings, by lowering the costs and improving the adaptability to wide areas of applications and domains.
- Expanding usage of robots in the medical sector, wearable robots and exoskeletons

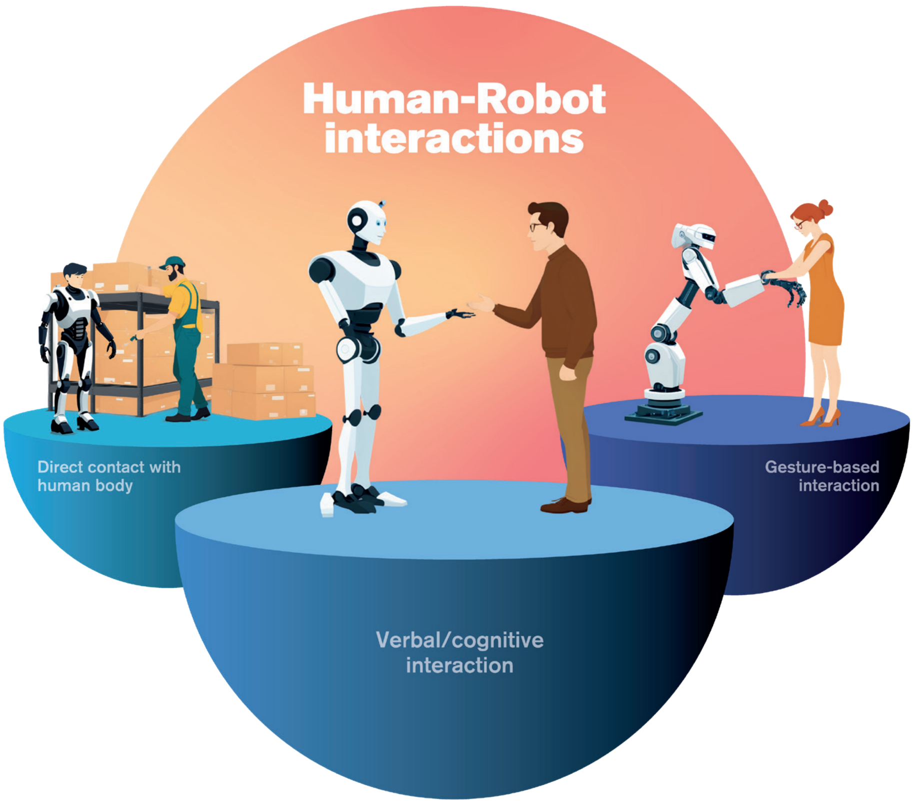

Fig. 5 Different types of human-robot interaction

- Remote intuitive interaction with robots and drones for the inspection of remote locations and for search and rescue operations in collaboration with humans, mobile robots capable of navigating in crowded spaces, such as hospitals, airports, restaurants, and robots for social companionship.

- Enabling the most challenging application, namely inside homes, the most unstructured and unpredictable environments.

collaboration and interaction with humans—the future will see robotic technologies becoming increasingly integrated with human beings. The integration of artificial and biological components will occur at all levels: physical connections between human and robotic bodies (such as prostheses, exoskeletons, and supernumerary limbs), as well as direct neural-synthetic signal interfaces in both efferent and afferent directions. These two intelligent systems—humans with their biological intelligence and robots with AI—

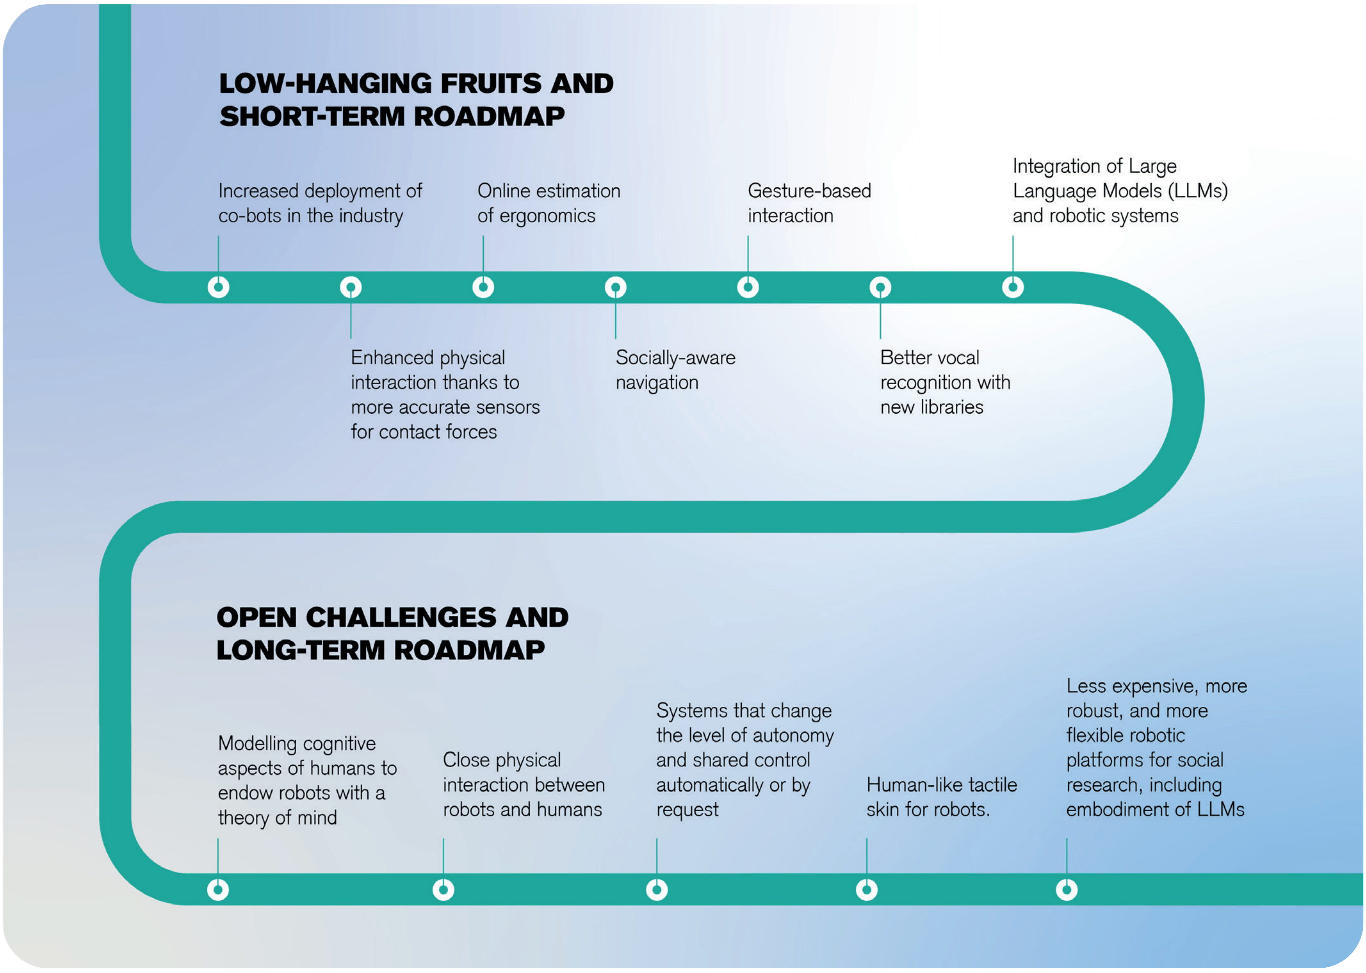


Figure 6: Short-term and long-term challenges in design and deployment of robots meant to interact with humans.

## HUMAN-ROBOT INTEGRATION: BIONICS AND BIOHYBRID ROBOTICS

If robots until the early 2000s were predominantly of the classical industrial type—accurate, rigid, and heavy—and in the first decades of the new century, cobots took center stage—robots designed for safe

will attain augmented capabilities unattainable by either system alone.

A paradigm for these future developments emerges from progress in rehabilitation robotics, where a new generation of prostheses, exoskeletons, and supernumerary limbs was developed—directly interfacing with the human body and neural systems to help patients perceive artificial parts as natural extensions of their own bodies. Cutting-edge technologies for surface, intramuscular, and intraneural high-density signal pickup and stimulation, along with cortical interfaces equipped with implantable devices containing thousands of neural sensors the size of a hair, are unlocking new possibilities for unprecedented advancements. Meanwhile, artificial musculoskeletal structures designed in alignment with ecological models of physical interaction—mirroring the environments our brains are naturally trained to navigate—are fostering the development of sensory-motor contingencies that closely resemble those found in nature.

It can be envisioned that the near future will see this paradigm pushed even further. As robots become increasingly capable of autonomous learning and replacing humans in repetitive tasks, a crucial frontier emerges in high-responsibility, one-off activities requiring human expertise and adaptability. In unknown and extremely unpredictable environments (such as our households), where neither humans nor robots alone can perform adequately, a new paradigm of human-robot symbiosis must be established. Moving far beyond traditional teleoperation, this vision of avatars as "whole-body prostheses" aims for seamless integration where humans and machines form unified symbionts, amplifying each other's cognitive and physical capabilities. Achieving this requires groundbreaking advancements not only in robotics and neural interfaces but also in neuroscience and philosophy, to unravel the deep connection between mind, body, and environment. Understanding the foundational elements of a shared cognitive and perceptual "language" between human and machine is essential for creating an immersive and empowering symbiotic relationship. Moreover, this shift raises new psychophysiological and ethical questions, as individuals experience out-of-body interactions through robotic embodiments, potentially redefining human identity and agency. This convergence of technological and humanistic inquiry gives rise to new scientific questions dedicated to exploring the profound implications of human-machine coalescence, with an immense potential for the development of new personal robotics assistants.

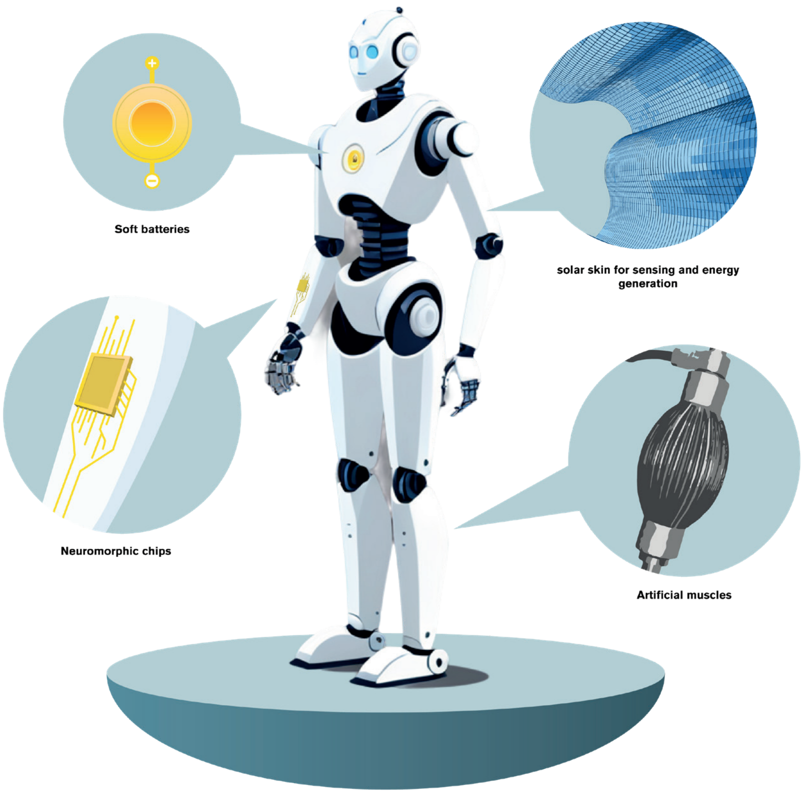


Figure 7: Humanoids and other platforms will benefit from the integration of new core robotic technologies to be developed over the next decades, including neuromorphic chips, tactile and solar skins, artificial muscles, soft batteries for energy storage.

## ROBOT BODYWARE

Several advancements in **core robotic technologies,** i.e. materials, sensors, actuators, energy storage, computing devices, are needed to obtain and turn into commercial products robots capable of interacting with unfamiliar, unpredictable and even harsh environments and collaborating safely and effectively with humans. In advancing these technologies the incorporation of new embedded functionalities or properties into the body/structure of the robot should be a guiding principle, along with the environmental sustainability of employed materials. The challenge is to design new classes of robotic components and systems (e.g., actuators/motors, sensors, structural parts, energy storage/management, etc.) that feature embedded characteristics such as resistance to challenging environments (space, underwater, radioactive, etc.), self-healing capabilities, biodegradability, and the capacity to recover and/or harvest energy from renewable sources. This requires careful consideration of the interplay between hardware and software, to find new strategies to co-design robot control and morphology. This will likely leverage on novel manufacturing and material processing technologies based on additive manufacturing and novel printing technologies.

For core robotic technologies to be successfully commercialized they cannot rely on expensive and hard-to-manufacture components, they must be robust, scalable, and affordable. Robustness ensures reliability and durability in various environments, scalability allows for adaptation to different robot embodiments and applications to respond to varying market demands, and affordability so that mature components can be integrated in a cost-efficient manufacturing process. These factors collectively ensure that the robot systems built from these technologies provide a return on investment (ROI).

Europe needs to engage in the very ambitious and long-lasting race for General Humanoid Robots, started today mainly in the US and China. Several cutting-edge prototypes of humanoids on wheels and on legs exist in Europe today, still being on par or even surpassing the competition. However, larger consortia between research and industry, including suppliers of components, need to be built quickly to keep up with the increased pace in international humanoids robotics. While torque-controlled, compliant robots developed in Europe have set the standard in interactive robotics over the last two decades, we need to look into new technologies, such high-efficiency actuators and combine them with new concepts of mechanical and electrical energy recuperation and storage, if humanoids robots are to do hard work continuously for

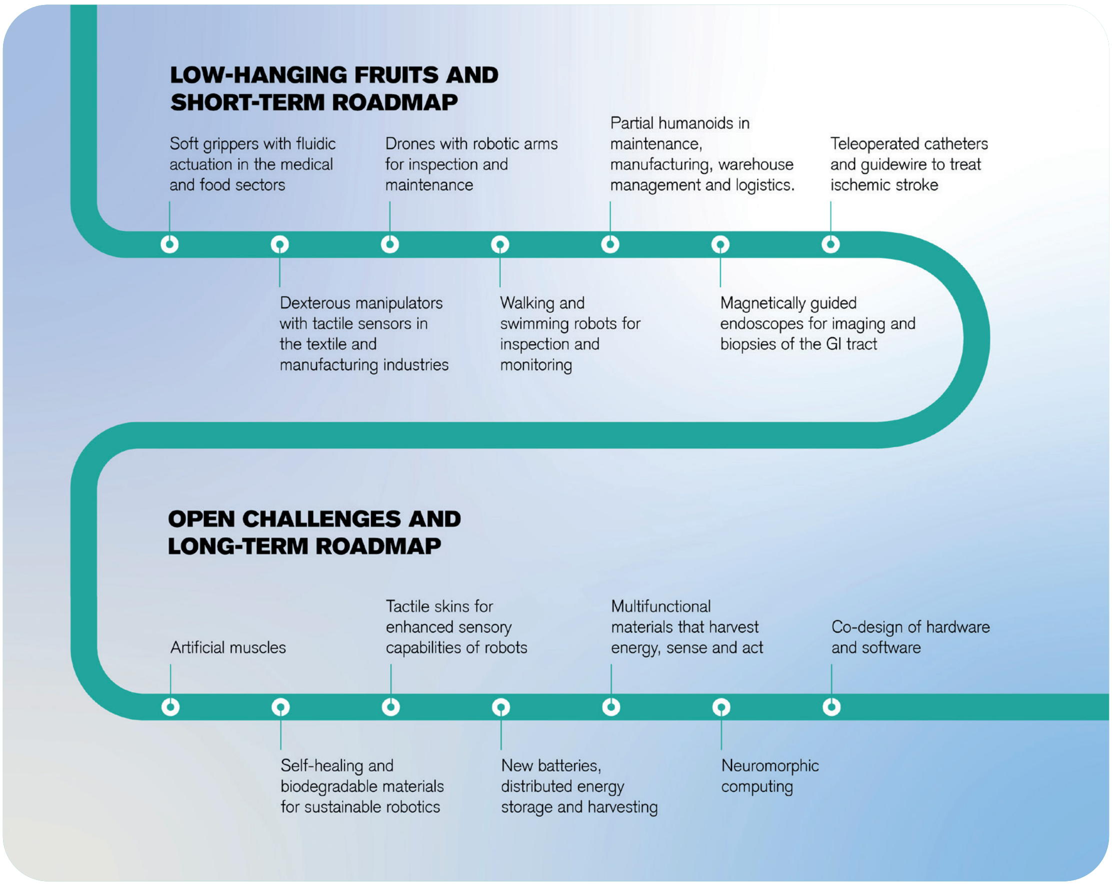


*Figure 8 Summary of short-term and long-term research goals for new robotic technologies and applications.*

many hours. Humanoids will need to be fail safe, i.e. legged robots need to fall over only very rarely and then survive the fall and not harm people in any case - a huge challenge towards reliability and safety of the entire system. Moreover, while most humanoids set the focus on legs and walking today, Europe's expertise in hands needs to be leveraged and extended, as manipulation is a challenge at least as big as locomotion and breakthroughs are still to come in this field.

**More versatile grippers** to tackle grasping and manipulation of delicate, deformable objects such as fruits, vegetables, garments can be obtained through novel designs, by co-designing grippers and their software, and with the addition of 3D tactile sensors. Remote control stations, supported by 5G/6G high bandwidth communication, can allow operators to control a robotic avatar over any distance with feedback of interaction forces. Steps can be taken towards new modular components with standard protocols that autoconfigure as part of a bigger complex robotics system like it is now the case for computer add-ons. Flying platforms can be equipped with robotic arms and grippers for aerial manipulation, with increased force and enhanced manipulation capabilities, to be used for inspection and maintenance of bridges, power lines, and high-rise buildings. Swimming and amphibious robots, which today are less advanced than legged and winged ones, can be improved and be applied for inspection of fish farms, submarine cables, and underwater platforms. Moreover, morphing robots will adapt to different environments and tasks. A breakthrough towards real-world applications in legged humanoids can be expected.

If robots are to approximate the versatility of living beings in facing tasks and environments, they need to replace at least some of their current electromechanics actuators with artificial muscles that have sufficient power output, are modular and self-healing. **Biohybrid robots** composed of a symbiosis between living tissues and cells (e.g. muscles) and artificial materials (e.g. prosthetic hand and exoskeletons) are already appearing today, and will be growing fast as 3D bioprinting techniques become more and more established. There will also be a need for flexible, robust, self-healing tactile skins, with high spatial resolution and force direction sensing to allow effective interaction with objects as well as safe interaction with humans. Edge AI systems to process tactile sensing data locally at the point of contact will help reduce the quantity of data to be processed by the central processing unit.

Current trends bet on **soft motors and sensors,** made of deformable, at times biodegradable materials and flexible electronics. The most mature technologies in soft robotics currently include fluidic and pneumatic actuation, electroactive polymers, shape memory alloys/polymers and the newly introduced electrostatic electro-hydraulic systems. Fluidic actuated soft robots can be used in simple grippers that can be applied in biomedical fields, for example for endoscopes, and in specific industries such as food and agriculture. Yet, current soft robots fall short of providing the strength and speed requirements to control full body robots capable of multi-purpose tasks. Also, design tools, material models and production technologies still require improvement to possibly reach higher TRL and subsequent transition from lab to market. The future likely lies in either designs that combine a rigid or semi-rigid skeleton with these new classes of soft actuators and materials, or any combination thereof.

**New sensors** that can be integrated into robots include ultra-low latency vision sensors, ultra-wide band localization, RF-sensing technologies, 3D force sensors, proximity sensing, sensors for human physiology and wearable sensors. Neuromorphic computing of spiking, event driven signals is expected to overcome classic computing architectures and communication theory limits, to increase performance and energy efficiency of robot electronics under the continuously increasing demand for computation and communication power.

**Improved energy efficiency** of onboard computation is needed to increase autonomy and operational life. For humanoids in particular this may require having distributed power generation along the body of the robot, instead of a backpack, which in turn will require novel materials science contribution. Ultimately, improved energy efficiency can be achieved through morphological computation, i.e. designing the robot soft body and generally the intrinsic body dynamics so that its shape and modes of deformation constrain its movement, reducing the need for computation and control of its states.

Some of the bodyware advances which will enable robots of the future are:

- New actuators with a high-power output, that are modular, redundant, efficient, capable of energy recovery, and self-healing

- Decentralized control to delegate portions of processing to subparts of the robot to ease modularity and speed up real-time reflex-like computation;

- Segmenting processors into smaller, specialized units, chiplet technology allows for more efficient and cost-effective designs; integration of AI accelerators and memory, within a single package, enhancing performance and reducing latency;

- Integration of batteries into the mechanical structure, possibly in a decentralized and distributed manner; design of soft batteries, with self-healing properties, and energy harvesting capabilities;

- Multi-faceted fault tolerance design that offer redundancy in actuation and sensing to enable robots to degrade gracefully in performance despite problems like a missing limb, electronic malfunction, or software error.

- Co-development of AI & microelectronic architecture to respect objectives of frugality, autonomy (embedded ad hoc architectures) as well as confidentiality.

- Technologies to record and deliver the signals from and to the human body, computational approaches to process them and interpret human intention, new haptic displays for conveying sensory feedback in a natural and intuitive way.

## ROBOTS IN THE REAL-WORLD

The manifold everyday applications for the next generation of robots require predictable control and explainable behavior which can be guaranteed if robots are endowed with reasoning abilities to encode and use semantic knowledge, to make inferences about the consequences of their actions and to make decisions. Significant progress in this direction has been made in AI-empowered robotics, with innovation continuing wherever data and/or models are available.

Unfortunately, some of the emerging societal challenges are exactly such that neither models nor data are available. Consider for example the most promising fields of application of Robotics:

- **Healthcare and Aging Population Assistance** (elderly care, including physical support, companionship, remote monitoring; exoskeletons to support weaker/aging workers in physically demanding tasks; autonomous disinfection robots and systems for healthcare logistics);

- **Medical and Surgical Care:** overcoming the “Robotic divide” that prevents most of the world population from the benefits of robot-assisted care: e.g. making medical and surgical robots affordable and accessible to everybody, everywhere.

- **Logistics** (delivery robots and drones for urban and rural applications, warehouse automation and inventory management, port and freight automation);

- **Agriculture and Food Security** (precision farming for seeding, planting, and harvesting to address labor shortage, monitoring crop health, irrigation, pest control, sustainable aquaculture and livestock management);

- **Construction and Infrastructure** (building and maintaining roads, bridges, and pipelines; autonomous 3D printing robots for sustainable construction, inspection and maintenance for aging infrastructure to improve safety);

- **Environmental Sustainability and Circular Economy** (waste sorting and recycling, renewable energy management (e.g., wind turbine and solar panel cleaning), monitoring and protecting biodiversity and natural habitats;

- **Civil security, Disaster Response and Climate Resilience** (Search-and-rescue robots for natural and industrial disaster scenarios, post-disaster infrastructure assessment and repair, firefighting, flood response, nuclear disaster, climate change monitoring and pollution containment);

- **Security and Defense** (law enforcement, border patrol, surveillance, drones for security monitoring and disaster relief, counter-terrorism, bomb disposal and reconnaissance);

- **Space Exploration and Industrialization** (Lunar and Martian exploration, autonomous rovers, habitat construction, maintenance of satellites and space stations mining and resource extraction for sustainable space industrialization);

- **Underwater Exploration and Industrialization:** While surface resources are finite and increasingly depleted, undersea resources remain largely untapped because of technological and environmental challenges. Polymetallic nodules are rich in precious metals and minerals, seafloor massive sulfides, rare earth elements deposits, methane hydrates, oil and gas and biotic resources from which we will get large part of our new pharmaceuticals.

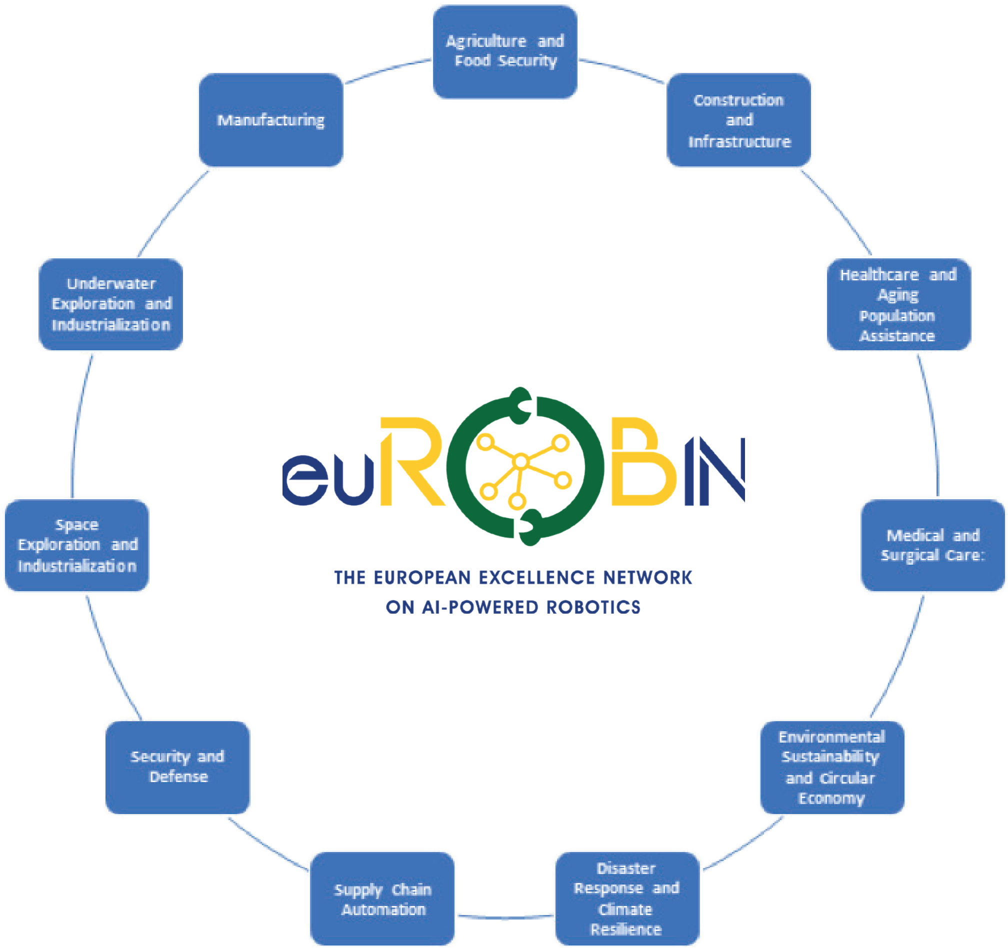


*Figure 9: Robotics and AI impact.*

- **Industrial production and Manufacturing:** while this field is the traditional playground for roboticists, there are still plenty of processes that could be improved by innovation in robotics, from logistics to lean manufacturing, from workers' ergonomics (physical and cognitive) to critical operations in harmful environments for humans, from new sectors like pharma and the food industry to artisan-based settings (where robotics is fundamental to support the digital revolutions)

These seemingly disparate applications really have much in common. Indeed, in most of these examples, the environment is totally or largely unpredictable, and neither mathematical models nor data are currently available to learn from. Novel strategies are hence needed to tackle such challenges, and new abilities are essential:

**1. Navigate uncertainty:** Develop systems that reliably function in complex, unpredictable and changing environments, without extensive historical data;

**2. Reduce data dependency:** Create adaptive algorithms capable of functioning with minimal data and generalizing from limited examples;

**3. Active Perception:** plan actions to look for missing data, by dynamically trading off goal attainment with information gathering, privileging e.g. innovation in data collection until enough are available to go for the task.

These applications share also the essential need of being "provably" safe, i.e. they must be tested and validated in all application domains. This is an extremely challenging task that is currently solved, e.g. for non-autonomous medical robotics, by validating the system for each specific anatomic district. A general purpose robot must handle each environment and its variants in a safe manner.

Any usage of robots in direct interaction with humans requires finding an intelligent way to combine model-based AI systems with deep learning algorithms, to mitigate potential risks such as misinterpretation. This requires defining in which situations misinterpretation can be accepted, because it poses no safety issues, and situations where we need instead that the machine really understands what happens, to assess it correctly, to avoid dangerous consequences.

Many future applications – from autonomous vehicles to prostheses and exoskeletons in the medical field – imply shared-control systems where humans delegate part of the decision-making and control functions to artificial agents. This creates the additional challenge of how to ascribe responsibility for failures and potential damage. A clear **regulatory and ethical framework** is needed, one with human needs and values firmly at its centre.

There is a fundamental need for **interdisciplinarity** at all levels from research to product design and deployment: Robotics cannot be designed by engineers and manufacturers alone anymore. The design and deployment of robots intended for direct human interaction and integration must be guided at every stage by expertise in psychology, ethics, law, and economics in order to ensure acceptability and economic viability in the market.

European institutions should weigh in with meaningful regulations to enforce the principle of human-centered robotics, as they have already done concerning the use and exploitation of personal data and the deployment of AI systems.

# CLOSING WORDS

Robotics has the potential to contribute to many of the United Nations' industry and environment-related Sustainable Development Goals. To achieve this vision, sustained investments in fundamental research and technology transfer, interdisciplinary collaboration and close interaction between academia and industry are required. Europe has a vantage position in robotics, and it will be crucial that it retains and strengthens the whole value chain, from design to manufacturing and deployment, with a balanced approach to AI development that requires investing not only in algorithms, but also in the ecosystems in which they will operate and in the underlying technologies.

Deployment of AI-powered robotics at large has become a more tangible target, possibly foreseeable in the next decade. Major hurdles on the road include ensuring understandability and controllability for safe deployment and usage and achieving scalable, cost-effective solutions to support autonomy and resilience. Further research and technological development is needed to improve system performance and fully exploit their potential.

Technological development must be accompanied by research on the social, psychological and legal dimensions of the relationship between humans and robots, to understand how humans can develop trust while avoiding excessive trust in robots, and how attitudes to robots change in time and across different cultures. This will ensure that future advancements in AI-powered robotics work in the interest of sustainable development, equality and social justice. Most importantly, sustainability should be included in developing robot technologies, considering life cycle extension, circularity, resource conservation and usage, ethics, and environmental justice to have a positive impact on the UN SDGs.

Any disruptive technology has an economic impact only if and when it is properly commercialized. In this respect Europe must strengthen its technology transfer programs, such as the EIC Transition and Acceleration programs, to support the transition from laboratory prototypes to commercial products, and foster the creation of more innovation ecosystems, capable of addressing the needs of deep tech entrepreneurs.

# ROBOTS AND SUSTAINABLE DEVELOPMENT

The next generation of AI-powered robots can help tackle key challenges faced by our societies:

An aging population: the need for assistance to the elderly and disabled in homes, or the need for physical and cognitive rehabilitation after incidents and disease, will greatly increase in the next decades, with simply not enough human caregivers.

Humanitarian responses during natural and man-made disasters that are predicted to become more frequent because of global warming, pollution and international crises. Robots will be increasingly needed for search and rescue, or for environmental remediation and decommissioning of industrial sites, including nuclear plants, and inspection of infrastructures after the disaster. As of today, robot-assisted surgical care covers less than 1% of all surgical interventions worldwide, depriving the vast majority of the population of the benefits of robot-assisted procedures. Affordable and accessible robot-assisted surgical and medical care is already scarce in Europe and will be only worse in the future.

The transition towards sustainable growth and a circular economy: robots can contribute to economic growth by increasing productivity in sectors that have not been automated so far, such as the textile or food industry, high-mix low-volume manufacturing, and maintenance of the European industrial and civil aging infrastructures. At the same time they can address the circular economy's increasing need to sort, recycle, and recover products and materials and keep them in the production cycle, including the handling of electronic components, batteries and toxic materials that should not be performed by humans.

Climate change mitigation: robots for environmental monitoring can contribute to a more precise assessment of the effects of climate change.

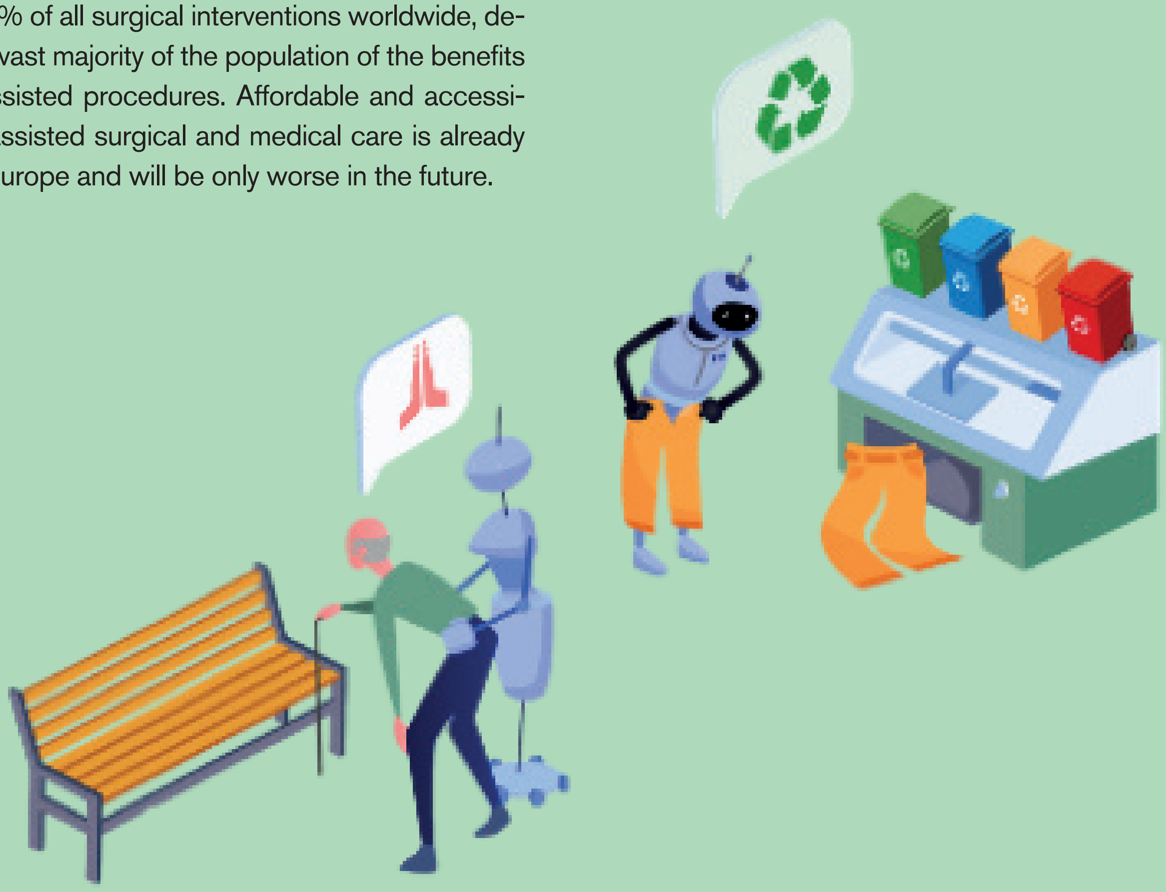

# AI AND ROBOTICS


From automatic translation to image recognition, from systems mastering complex board games to ChatGPT and the other language models, deep learning has achieved a lot in the last decade, and expectations on future developments of AI could not be higher.

Robotics has greatly benefitted from advancements in machine learning: for example, locomotion in legged robots has advanced greatly thanks to reinforcement learning, that allows to define a high-level target such as the speed of locomotion or a destination without a full mathematical description of the problem. Thanks to the advancements in deep learning, driverless cars are being tested as commercial service in major cities. Robot simulators have advanced thanks to deep reinforcement learning, which allows exploring policies with different environmental conditions in a reasonable amount of time before trying them on the actual robot.

But unlike language models and image recognition algorithms that only deal with bits, embodied AI poses specific challenges. Robots cannot rely on huge datasets that can be digested in relatively short times. Datasets themselves based on physical interactions (locomotion on different terrains or grasping of various objects) are simply not available and cannot be quickly assembled: having robots execute tasks in the real world takes time, and risks damaging the robot or its environment when attempts go wrong. A training dataset for flying robots, for example, would need to be impossibly huge, since drones can fly at vastly different altitudes and tilting positions with respect to the ground. The use of simulators is of great help, but for many tasks sim-to-real transfer is still a challenge.

Another crucial difference with non-embodied AI is that robots often perform safety-critical applications, and safety agencies would not approve a robot powered by an AI without enough transparency on when and why it may fail.

For this reason, AI-powered robotics will most likely include deep learning in combination with models that incorporate fundamental knowledge about the world and use it to guide and constrain the use of learned policies. Similarly, the development of AI-powered robots must go hand-in-hand with the development of testing tools that must be as advanced as the robots themselves.

Ultimately, because no data set or simulation can live up to the complexity of real-world physical interaction, robots will require lifelong learning and transferability of knowledge across tasks, across robot bodies and across environments, as well as between humans and robots. Research will need to focus on understanding what to transfer (identifying relevant knowledge about environments, objects, and tasks constraints); how to transfer (formalizing prior knowledge on robot bodies and sensors, kinematics and dynamics, and for a given task/ environment/body find feasible sets of motor commands); and when to transfer (learning to recognize similarities across environment, objects, and tasks constraints).

# HUMAN-CENTERED ROBOTICS

Future robots be they humanoids, drones, legged robots, manipulators, or entirely new soft robots are expected to operate in much closer contact with humans, collaborating and interacting with them in homes and offices as well as in public spaces. Ultimately, the vision of AI-powered robotics is to enable humans and robots to share spaces and tasks, deciding and acting together, while preserving humans' privacy and autonomy. This creates several new challenges. On the technical side, we need to devise and build cognitive and interactive abilities that allow pertinent, transparent, and legible behaviours in robots, a necessary premise to ensure that they can be trusted to work in collaboration with humans. On the safety side, we need to evolve current safety standards so that they account for the use of robots not only in private, controlled spaces but also in public, crowded ones: robots must be able to account for the heterogeneity of pedestrians, the dynamics of crowds, for social norms, and for real people's disorderly and at times mischievous behavior.

Many future use cases – from autonomous vehicles to prostheses and exoskeletons in the medical field – imply shared-control systems where humans delegate part of the decision-making and control functions to artificial agents. This creates the additional challenge of how to ascribe responsibility for failures and potential damage. A clear regulatory and ethical framework is needed, one with human needs and values firmly at its centre.

It is only through tight coordination with human-centered disciplines such as ethics, psychology, social sciences, that robotics can deal with the social, societal and ethical issues related to the use of autonomous machines in professional, public and domestic environments.

# HUMAN-ROBOT INTEGRATION FOR BODILY INTELLIGENCE

Robotic technologies are increasingly integrated with human beings. To describe the meaning of integration, consider two simple examples: robotic prostheses, such as hand prosthetics, and surgical robots controlled by surgeons. In both scenarios, human capabilities are enhanced: an amputee regains hand functionality, and a surgeon operates with higher dexterity through minimally invasive procedures. In both cases, the artificial system does not operate autonomously and requires an interface to integrate with the human user. In both cases, however, many services can be provided autonomously by the artificial part of the system. This is precisely our vision: the integration of two intelligent systems, humans, with their biological intelligence, and robots with AI to achieve augmented capabilities unattainable by humans alone.

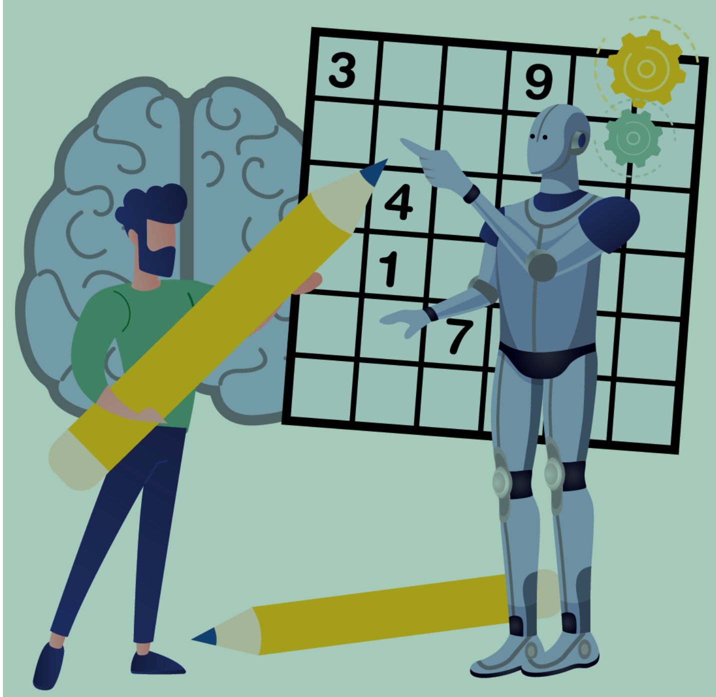


We are moving toward a world where robots and artificial intelligence will reach increasingly advanced performance levels. To avoid being left behind in shaping the future society, a society of humans, robots, and AI, we must synergize with artificial systems, ensuring humans retain a central role. A crucial role, indeed, akin to the "smart keys" of modern cars that deactivate the overall operative system of the car when the key is absent. When referring to augmentative robotics, we are referring to wearable robotics, supernumerary limbs and avatars, and to all complex integrated human-robot systems where humans always maintain control and awareness of the task.

Examples of integration currently yielding significant results are primarily in healthcare applications, such as robotic prostheses, supernumerary limbs, or surgical robots. In these cases, the presence of humans is non-negotiable, as the decisions to make are crucially related to people's health and safety.

While today's most impactful applications are medical, we must prepare for a future where these integrative technologies will be employed pervasively in other fields. For example, supernumerary limbs and avatars can be used in logistics for object manipulation up to space exploration, where the first few humans traveling to Mars will require additional arms to build habitable capsules. We want to prepare for an era where humans control multiple artificial and robotic systems simply as their extensions, just like their eyeglasses or their smartphones, but with an ability to act in the physical environment.

# ROBOTS IN THE WILD: TACKLING THE REAL WORLD CHALLENGE

Significant progress has been made in robotic and AI solutions, with innovation continuing wherever data and/or models are available. Unfortunately, some of society's emerging challenges are exactly such that neither models nor data are available.

What are the most urgent and promising fields of application or Robotics and AI and what do they have in common?

Healthcare and Aging Population Assistance, Agriculture and Food Security, Construction and Infrastructure, Environmental Sustainability and Circular Economy, Disaster Response and Climate Resilience, Supply Chain Automation, Security and Defense, Space Exploration and Industrialization Underwater Exploration and Industrialization: in common among these apparently disparate applications is that in most cases the environment is totally or largely unpredictable, and neither mathematical models nor data are available to learn from.

Novel strategies are needed to tackle these challenges. To overcome them, it is essential to:

**1. Navigate uncertainty:** Develop systems that reliably function in unpredictable and complex environments, even without extensive historical data.

**2. Reduce data dependency:** Create adaptive algorithms capable of functioning with minimal data or generalizing from limited examples.

**3. Active Perception:** plan actions by dynamically trading off goal attainment with information gathering, privileging innovation in data collection until enough are available to go for the task.

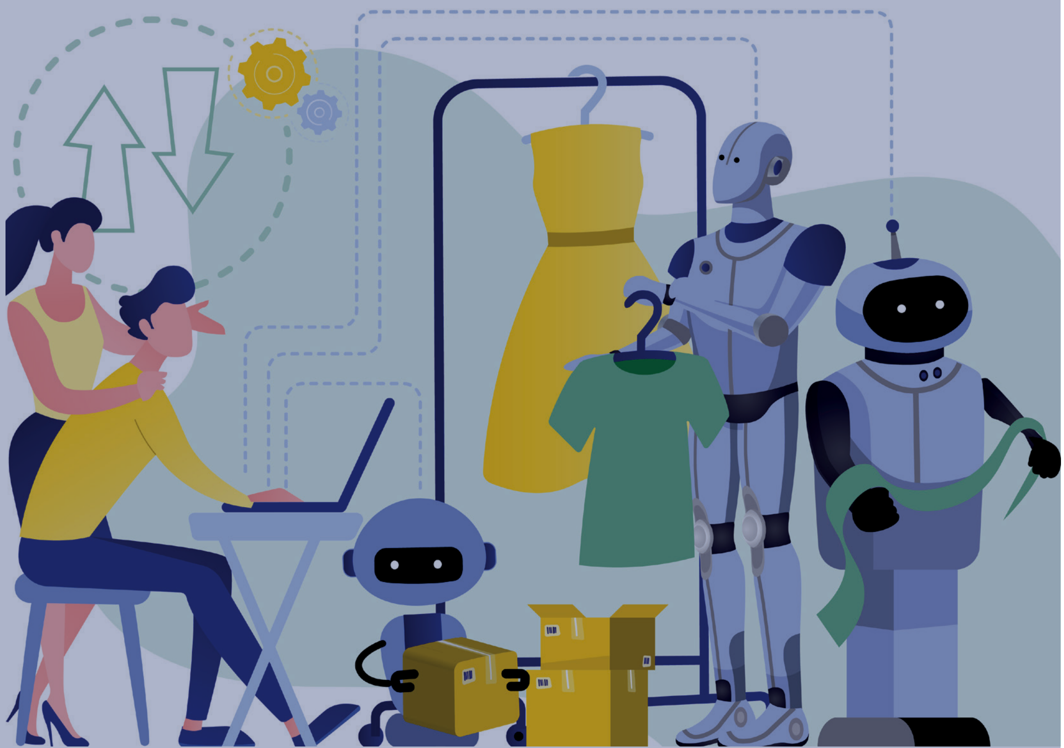

THE EUROPEAN EXCELLENCE NETWORK ON AI-POWERED ROBOTICS

**euROBIN website**
(with new EuroCore and SRA features)

**euROBIN Week 2024 -**
Video of the 1st euROBIN Coopetition @Nancy

@euROBIN_project

@euROBIN-project

eurobin-project.eu

This project has received funding from the European Union's Horizon Europe research and innovation programme under grant agreement No. 101070596

# ANNEXES

The following perspective articles expand the key themes of the Strategic Agenda, and are currently under review at Nature Machine Intelligence (1) and IEEE Robotics and Automation Magazine (2-4).

THE EUROPEAN EXCELLENCE NETWORK ON AI-POWERED ROBOTICS

This project has received funding from the European Union's Horizon Europe Framework Programme under grant agreement No 101070596.

# 1. Will AI alone solve robotics?

**Aude Billard**, Learning Algorithms and Systems Laboratory (LASA), École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland

**Alin Albu-Schaeffer**, Institute of Robotics and Mechatronics, DLR-German Aerospace Center/Department of Informatics, Technical University of Munich

**Michael Beetz**, Institute for Artificial Intelligence, Computer Science Department, University of Bremen

**Wolfram Burgard**, Department of Engineering, University of Technology Nuremberg, Germany

**Matei Ciocarlie**, Columbia University, New York City, USA

**Peter Corke**, Center for Robotics, Queensland University of Technology, Brisbane, Australia Matei Ciocarlie, Columbia University, New York City, USA

**Ravinder Dahiya**, Northeastern University, Boston, USA

**Danica Kragic**, School of Computer Science and Communication, Royal Institute of Technology, Stockholm, Sweden

**Ken Goldberg**, University of California, Berkeley, USA

**Yukie Nagai**, International Research Center for Neurointelligence, The University of Tokyo Davide Scaramuzza, Robotics and Perception Group, University of Zurich, Switzerland

**Davide Scaramuzza**, Robotics and Perception Group, University of Zurich, Switzerland

Deep learning, large-language models, and other AI technologies have gone from one breakthrough to the other. As a result, we are witnessing growing excitement in robotics at the prospect of leveraging the potential of AI to tackle some of the outstanding barriers to the full deployment of robots in our daily lives. However, action and sensing in the physical world pose greater and different challenges than analysing data in isolation. As the development and application of AI in robotic products advances, it is important to reflect on which technologies, among the vast array of network architectures and learning models now available in the AI field, are most likely to be successfully applied to robots; how they can be adapted to specific robot designs, tasks, environments; which challenges must be overcome. This article offers an assessment of what AI for robotics has achieved since the 1990s and proposes a short- and medium-term research roadmap listing challenges and promises. These range from keeping up-to-date large datasets, representatives of a diversity of tasks robots may have to perform, and of environments they may encounter, to designing AI algorithms tailored specifically to robotics problems but generic enough to apply to a wide range of applications and transfer easily to a variety of robotic platforms. We close on what we view as the primary long term challenges, that is, to design robots capable of lifelong learning, while guaranteeing safe deployment and usage, and sustainable computational costs.

# 1. INTRODUCTION

The last decade has witnessed impressive advancements in the development and practical application of Artificial Intelligence (AI) technologies, in particular for systems based on Deep Learning (DL) over multi-layer artificial neural networks (ANNs). Though ANNs are not recent concepts, several factors have contributed to a fast-paced acceleration in their performance and scalability. On one side, computing platforms, such as Graphical Processing Units (GPUs), have become available, offering increased computational power and allowing to create "deeper" networks (i.e. with more hidden layers). On the other side, the exponential growth of multimodal, digital information available on the Internet has made vast amounts of data easily available for the creation of training and test datasets.

The first demonstration of the potential of these technologies came in the early 2010s, when deep networks started overcoming previous systems in visual recognition challenges[1]. Since then, there have been important applications of these systems on several different computational tasks.

Great expectations currently surround the applications of new AI systems (including ANNs, DL and LLMs) to robotics. Once again, this is not a novel concept, because learning algorithms have been used to control robots for decades. But there is hope that the current fast-paced scaling-up of AI's performances may translate into a similar scaling-up of robotic capabilities and help solve some long standing challenges that have so far limited robots' autonomy in challenging environments or their capability to interact effectively and safely with humans. For example, the classic control and state estimation methods for robots, that were developed for industrial applications in controlled environments, struggle to adapt to the high complexity and intrinsic unpredictability of outdoor natural environments, or even to the diversity of objects that can be encountered in a typical home. It is tempting to expect that advancements on these problems will mirror what happened for Go – a boardgame that was famously impossible for classic computer programs to master mathematically. Deep Learning came and vastly surpassed human abilities, albeit after playing billions of games with itself.

However, we cannot expect that what worked so well for perfect information games, which are purely data-based, software-level tasks, such as image recognition or text generation, can be applied with the same success to sensing, planning, control, and navigation for physical machines. Action and sensing in the physical world pose greater and different challenges than playing games: the state space is bigger, training data are not so easily available and cannot be easily generated, and safety and reliability requirements are higher. It is then paramount to identify which technologies, among the vast array of network architectures and learning models now available in the AI field, can be successfully applied to robots and which cannot; how they can be adapted to specific robot designs, tasks, environments; which challenges must be overcome.

# 2. A BRIEF HISTORICAL REVIEW

Allowing robots to operate autonomously in novel situations and to approximate the dexterity and agility of living organisms have been key challenges for robotics since at least the 1960s[2 3 4]. For several decades, robotics researchers have been experimenting with neural networks and machine learning as a potential solution to those challenges, and there is now a sizable literature on how to leverage these techniques to tackle robotics problems that had previously proven hard to solve. These studies have provided insight into which styles of machine learning are most suitable for robots, and which tasks are more amenable to be learned rather than formally programmed.

Overall, two principal styles of machine learning have been employed in robotics since the 1990s. On one side there is a family of algorithms that allow robots to learn from expert data, typically provided by a human demonstrator who performs the target action while their movement is captured by visual or motion sensors. Called alternatively Programming by Demonstration, Learning from Demonstration (LfD) or Imitation Learning, this approach has proved applicable in tasks ranging from grasping to manipulation of complex objects[5 6 7]. LfD algorithms could produce impressive results, such as catching objects in flight or control complex flying manoeuvres[8 9], while relying on very small datasets. The main limitation of LfD has historically been the intrinsic need to have a human operator with a good knowledge of the task available for training the robot, often across many training sessions. To address these challenges, current efforts are directed to learning from non-experts or suboptimal demonstrations, or from large collections of human and robot actions[10 11 12]. Other approaches, such as active learning[13], one-shot and behavioral imitation[14 15] or behavioral cloning[16], have also been proposed as a way to improve the efficiency of LfD: these techniques allow the robot to query the expert for demonstrations only when required, to learn a complete behavior from a single demonstration, or to start by acquiring experience in a self-supervised fashion and then use this experience to develop a model which is then used to facilitate learning of particular task by observing an expert. All of these have been shown to require fewer post-demonstration environment interactions than other techniques.

The other type of learning algorithms enables robotic systems to learn through trial and error without a prior formalization of what constitutes the correct control policy. Best exemplified by reinforcement learning (RL)[17], this method typically relies extensively on computer simulations of the robots and its environment to create enough learning cycles and learn a robust enough policy before testing it on the actual robot. Use of RL in robotics was hindered, for a long time, by the exploration phase, which, if not properly bounded, can become too computationally and time intensive and its inability to easily scale to high dimensions. Recent advances leverage the increasing effectiveness of deep-learning and visually-realistic physics-based simulation, achieving notable success in applications such as locomotion for legged robots – both quadrupeds and humanoids – as well as flying robots[18 19 20]. These methods are limited in that training must be conducted initially in simulation, far from the complexity of the real world, and the transfer from sim-to-real remains an issue[21]. In addition, RL success depends on a good prior knowledge of how to define an effective reward metric and assess the robot's performance against it.

Some of these challenges can be resolved when using LfD and RL in combination to leverage the strength of both techniques while mitigating their limitations. LfD can be used, for example, to reduce the search space in RL by bootstrapping it with good examples[22], reducing training time of large models[23], or to infer the reward and the optimal control policy simultaneously, a technique known as Inverse RL[24].

## 3. POTENTIAL FOR NOVEL APPLICATIONS AND COMMERCIAL DEPLOYMENTS

Many advances initiated in academic research have found their way to commercial applications. AI powered robots that can pick and sort packages of various sizes are regularly deployed in e-commerce warehouses. Learning enables online adaptation in tasks like pick-and-place on assembly lines, which were once rigidly pre-programmed. Robots can now adjust trajectories if an object is misplaced, or its shape or weight is unexpected. Autonomous cars, which started in the early 2000s, are now commercially deployed – ranging from partial autonomy in most models currently on the market to pilots of full autonomy underway, in limited situations, in several cities.

While AI is now pervasive in all areas of robotics, an area of application of particular interest is the field of soft robotics, where the deformable, continuum nature of robot bodies and their complex interaction with environments makes the processing of sensor data, state estimation, and control particularly challenging. Soft robotics is regarded as one of the many promising areas in robotics. Its natural compliance may ease the usage of robots in areas requiring direct interaction with humans and address global issues through biodegradable solutions. AI may offer an alternative to traditional control methods that cannot be used readily to control soft robotics and process their complex and heterogeneous sensor data stream[25], thereby easing usage and deployment of this new technology. A notable example is the recent application of convolutional neural networks to interpret the wealth of data streaming from a soft glove's artificial skin, enabling real-time recognition and control of grasps on objects[26].

## 4. SHORT- AND MEDIUM-TERM CHALLENGES

Scientists have only begun to scratch the surface of the potential of RL, LfD, and other flavors of AI and machine learning for robots. We next point out a list of short-term and long-term challenges, by increasing level of complexity, all of which form the corpus of current ongoing research directions (see Figure 1).

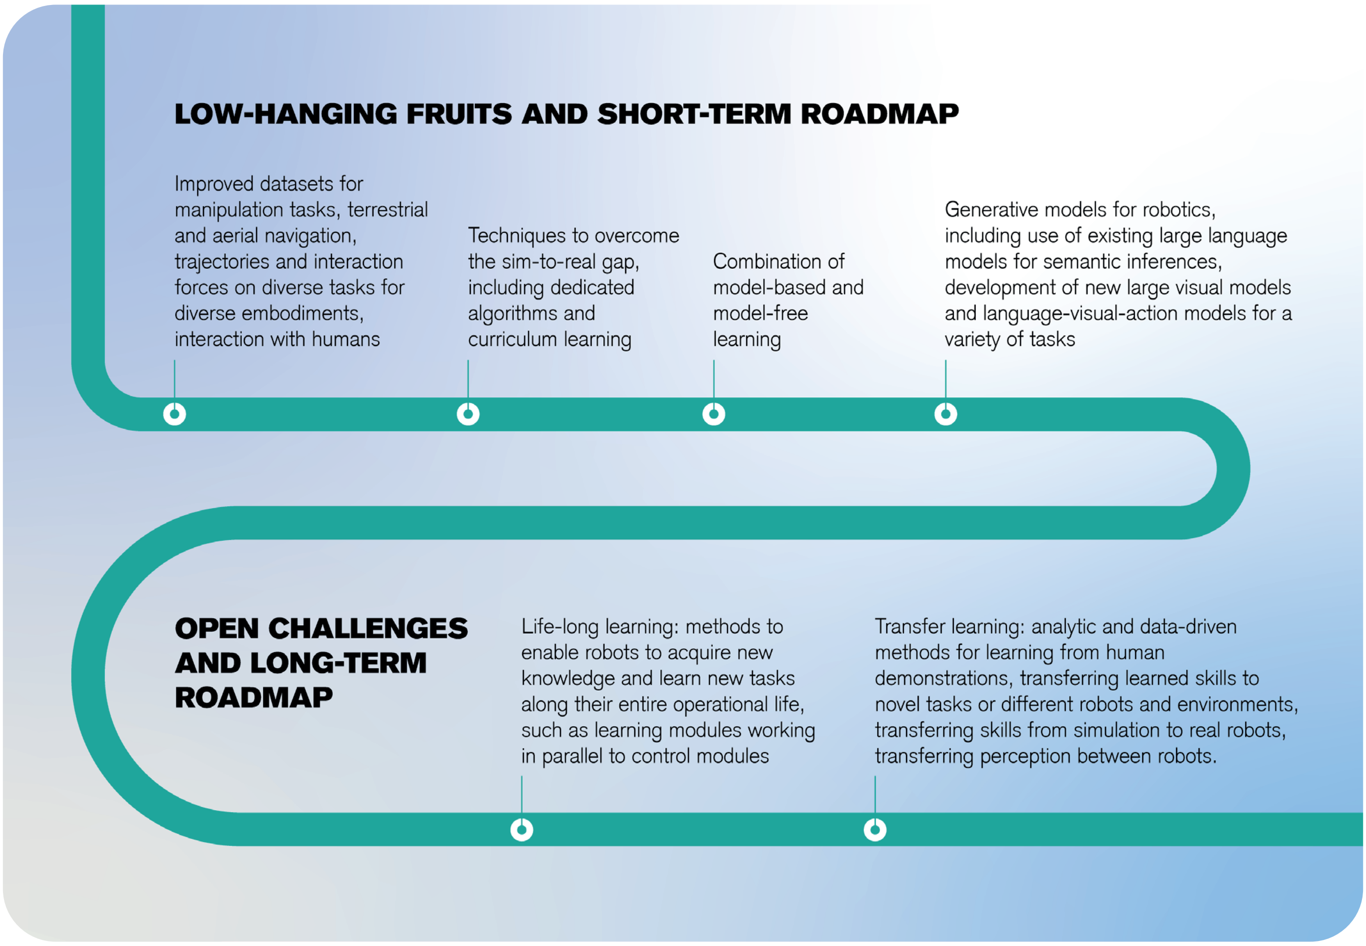


*FIGURE 1: Short-term and long-term challenges for further endorsement of AI in robotics, order by increasing level of complexity. Note that these challenges may not be overcome sequentially. Rather, research proceeds in parallel along many of these directions.*

**Creating and maintaining representative datasets.** An intrinsic limitation in robot learning as compared to other AI application domains is that there are no ready-to-use and easily available large datasets that can be used to train ANNs on sensing and control tasks, comparable to the vast repertoire of images and text that could be downloaded from the Internet and used to train image recognition or text generation algorithms. Generating *ex novo* enough iterations of a robotic task to train an ANN can be exceedingly costly and time consuming, or simply impossible. Too many robots would be destroyed during failed attempts at a task, and in some cases (such as autonomous flying robots) this would create risks for humans.

For some tasks, reference databases can indeed be created but require an organized and multi-centric effort. For example, in the case of visual imitation learning, attempts are being made at creating an analogue of ImageNet for grasping and manipulation, such as the Dexterity Network (Dex-Net) research project that develops code, datasets, and algorithms for generating parallel-jaw robot grasps and metrics of grasp robustness based on physics for thousands of 3D object models, with the aim of training machine learning-based methods to plan robot grasps. It supports researchers in finding robust grasps and training neural networks to generate a wealth of similar grasping strategies. The platform has allowed to learn deep policies to pick objects from a bin containing many unfamiliar objects at various orientations, the so-called “bin-picking” problem that has long been a benchmark challenge in the field[27].

Large datasets are also being created for terrestrial navigation tasks, thanks to cars now collecting large amounts of images routinely, from professional mapping services such as Google Maps to dashcams becoming increasingly common on private vehicles. These databases are typically available to companies on a proprietary basis, but if privacy and IP issues can be dealt with, it is foreseeable that some of them can become available to researchers. The challenge is bigger for aerial navigation, because of the many different perspectives from which a drone can observe the same scene, at vastly different altitudes and tilting orientations with respect to the ground.

Beyond visual data, robot learning needs datasets of robot actions in the form of trajectories and interaction force profiles associated with various tasks. Datasets on specific robot bodies and tasks do exist, but they are typically too narrow for large-scale machine learning. Combining datasets from diverse embodiments and on diverse robotic tasks can be a solution to reach the required scale. For example, an effort has recently been launched to combine several datasets on robotic manipulation, each one based on a specific robot body and skills and has provided a proof of concept that such a combined dataset could be used to train a policy for a given task more effectively than by using a dataset specific to that task[28].

Possibly the biggest challenges in terms of dataset creation are related to close interaction with humans, as the complexity and variability of both physical interactions and communication with humans and the need for enhanced safety guarantees currently prevents the rapid creation of datasets either through real experiments or in simulation. Ethical issues also put strict limits on what data can be collected and stored about human subjects and how they can be labelled, for example, by ensuring that subjects are not recognizable, that no sensitive information about them can be inferred from the data, or that images of a human subject cannot be reused in a different context, including being used for different training objectives than initially specified. An additional complication is that robots and humans perceive the world and interact with it in very different ways: while humans rely on multimodal information combining visual, acoustic, and haptic information, robots mostly rely on vision or on other bands of the electromagnetic spectrum, and while they can see more than humans do (including in low light or through obstacles) they remain incapable of analysing complex visual scenes.

**From simulation to reality**. Simulations offer a partial solution when it is not possible to create a large enough dataset. Several robotic simulators are available to the robotic community (examples include Algoryx, Bullet, Gazebo, Isaac Sim, MuJoCo, RoboDK,) and have been used for a long time to test and improve classic model-based control algorithms before applying them to real robots. The accuracy of their real-time physics-based simulation (the so-called physics engines) has greatly improved, also thanks to their commercial use in computer gaming. Reliable physics-based simulators can now, for example, simulate locomotion on complex terrains and manipulation on realistic objects in home environments, allowing the evaluation and selection of optimal controllers in simulation before being downloaded in the real robot. The use of simulators reduces the time needed for training, requiring only fine-tuning of parameters on the real robot. Randomized control policies generated by a neural network can be run in simulation over several thousands of iterations, generating a training set from which optimal policies can be learned and then transferred to the real world.

However, overcoming the sim-to-real gap, i.e. the discrepancy between the robot's performance in the real world and in the simulated environment, remains a challenge. This gap can be the result of multiple factors: the simulator's model can be exceedingly simplified with respect to the actual physical robot; the variability of the environmental conditions can be too large to be captured by a model; the physics simulator can fail to accurately capture the physics of the real world, especially when it comes to contact forces and deformable surfaces.

There are many techniques to overcome the sim-to-real gap. A small amount of data from the real world can be collected and used to increase the realism of the simulator[29]. Rapid-Motor Adaptation (RMA) has been successfully applied to achieve online real-time adaptation of quadruped locomotion to changing terrains, payloads, wear and tear[30]. "Curriculum learning", where the robot learns gradually more complex tasks in gradually more complex environments, has also been shown to improve the transfer of policies learned in simulation to the real world for legged robots' locomotion[31].

**Leveraging large generative models for robotics**. Much of the current excitement around AI focuses on generative AI, and specifically on Large Language Models. They are mostly based on the "transformer" deep learning model, which around 2017 emerged as an alternative to both recurrent and convolutional neural networks, allowing the speedup of learning (in particular of textual information) by processing information sequences in parallel[32]. By learning statistical relationships in text documents, these systems have achieved remarkable efficiency in generating text. Based on the same principle, they have been applied to diverse problems such as computer coding or computational chemistry.

LLMs are attractive for robotics on multiple levels. Existing LLMs can be adapted to support human robot interaction based on natural language, essentially making it easier to control a robot through written or verbal instructions, in any human language, and allowing them to respond to humans accordingly. Attempts are also being made at using LLMs in robot navigation in new and unfamiliar environments, to support semantic guesswork, essentially using their inferences[33].

Another family of generative models are language-vision models, that are trained on text/image pairs or annotated videos found on the Internet, and that can be used to generate synthetic images and videos from text prompts[34]. These models can also be applied to robotics, for example, to improve object recognition in manipulation and navigation tasks, and allow tasks to be specified in terms of what can be seen by the robot. A new generation of large visual models can be purposely built for robotics, trained not (or not exclusively) on text/image pairs from the Internet, but on navigation datasets such as those described in the previous section, produced by cameras during actual navigation in real environments. A first step could be learning to generate expectations on domestic spaces, i.e. using datasets of images of homes and offices or information from sensorised objects to generate reliable predictions on what a robot moving around such an environment may encounter. The same approach could then be extended to terrestrial and aerial navigation, creating models that can understand and contextualize visual information and incorporate a model of the robot's own physics and behavior to predict what it will see next.

The most recent developments in the field are language-vision-action models that add action to the equation. Examples of such models are being proposed, trained by fine-tuning vision-language models with both Internet-scale visual-language tasks and robotic trajectory data. By expressing the robot actions as text tokens and incorporating them into the training set together with natural language tokens, these models can learn to output robot actions like LLMs output text[35]. Initial results are encouraging, but the challenge of feeding such models with suitable datasets (see section 2.1), effectively mapping vision to action, and providing the system with the reasoning capability to correctly anticipate the consequences of its actions, will have to be a core research focus for several years. Another challenge is to verify the logic and feasibility of the plan generated by LLMs, an issue that is well addressed in logic-based planning[36].

**Prior knowledge and combining AI with control methods**. For physical robotics, incorporating prior knowledge on both robot and environment dynamics in combination with control methods with provable guarantees, is a more sensible way forward than a totally bottom-up, knowledge-agnostic approach to learning. In aerial robotics, for example, neither learning nor aerodynamics-based control alone can help solve the challenge of approximating the agility of birds' flight: coupling sensing and perception with the full body dynamic, allowing a drone to have instant reactions in flight and cancel perturbations, or on the contrary profit from the wind, efficiently combining flapping of wings and gliding (in the case of a winged drone) to save energy. These challenges will require a combination of learning for building improved aerodynamics models with control methods for guaranteeing flight stability.

Another reason for combining models and formal knowledge with machine learning is that a system only based on the latter would be prone to failures that can neither be predicted beforehand nor fully explained afterward, as exemplified by "hallucinations" observed in Large Language Models. Many current deep learning models are intrinsically non-explainable, a problem that becomes even more critical when AI is applied to robots. And because most future robots are expected to operate in safety critical scenarios such as autonomous navigation or close interaction with humans, no regulatory agency would approve their use unless their behavior can be predicted, and performance guarantees met – failures must be explained and corrected which is currently not feasible with model-free deep learning. This is a serious limitation to almost all applications where harm to humans is possible, for example, in the medical field, aeronautics, logistics and transportation, and domestic use.

AI-powered robots will need models of the actions that they are about to do, and these models must be explicitly represented in order to reason about the consequences. For example, a robot designed to work in a chemical lab, whose task is to pour chemicals into different containers, needs to know what happens when an acid is mixed with a base. Whenever a human comes into play, the robot needs an actual theory of mind modelling what the human may do and how the human might interpret the robot's task. This model can quickly become more complicated than the model of the robot itself[37].

Numerous efforts are directed to merge control theory and machine learning, proving that this can ultimately speed up learning, increase the robustness of the learned model, and enhance its safety[38,39]. For example, a standard machine learning algorithm optimization can be modified to encompass penalties for violations of Lyapunov stability, or bounded constraints to guarantee estimated plausible values for physical quantities such as stiffness and mass[40]. In a similar vein, training of deep RL can be guaranteed to generate stable trajectories[41], or be enhanced by incorporating reference motions generated through control model, covering a broad range of velocities and gaits[42], serving as targets for the RL policy to imitate. Control theory and deep learning have also been combined to optimize grasps, using DL to find an initial policy that is then refined with model-based algorithms, thus sizeably speeding up computing[43].

# 5. LONG-TERM CHALLENGES

The most exciting, but also most challenging, long-term promise of AI for robotics is to enable robots to continuously acquire new knowledge, a dream dating back the 90's[44]. It requires three ingredients, which we discuss next.

**Life-long learning**: If the goal of robot learning is to approximate the way living organisms – humans included – learn tasks, then future robots will need to be able to acquire new knowledge and learn new tasks along their entire operational life, instead of relying on an initial training dataset that could never prepare them for the complexity and variability of the real world.

Endowing robots with the ability to learn continuously poses huge technical and regulatory challenges. Lifelong learning requires new paradigms based on incremental learning and is able to convert input output learning to structured knowledge, combining the power of learning with the paradigms of expert systems. It requires a learning module working around the clock on the robot in parallel to the control module enacting the policies that were already validated. It brings along difficult questions, such as: how do we get some assurance about the performance of the system? How can we test the system, provided we can't know in advance the situations it will encounter and how it will learn from them? How do we select the things the robot can forget to make room for learning new things? How do we make sure that whenever it learns something new, even minimal, it has not forgotten how to do something important that it could do yesterday? These problems will need to be investigated in close collaboration with neuroscientists and developmental psychologists, to understand how humans achieve continuous and diverse cognitive development transitioning from one task to another, how this mechanism can be reproduced in neural networks, and how they can be implemented in robots. These problems will also translate into major regulatory issues: how to check that an evolving system maintains the safety and reliability standards requested for market certification as its capabilities change with new learning?

Possibly the main challenge for life-long robot learning will be to be able to scale up the current learning methods. Many robots will not stay the same for their whole operational life. After five or eight years of operation, a robot may have to mount a different gripper, or a different motor. The objects it has to manipulate and the environment in which it operates may also have changed. When that happens, the acquired knowledge that allows the robots to pick up and manage different objects may not automatically transfer to a slightly modified platform. But we currently lack good algorithms to transfer automatically, without retraining or human intervention, across even small changes in the embodiment.

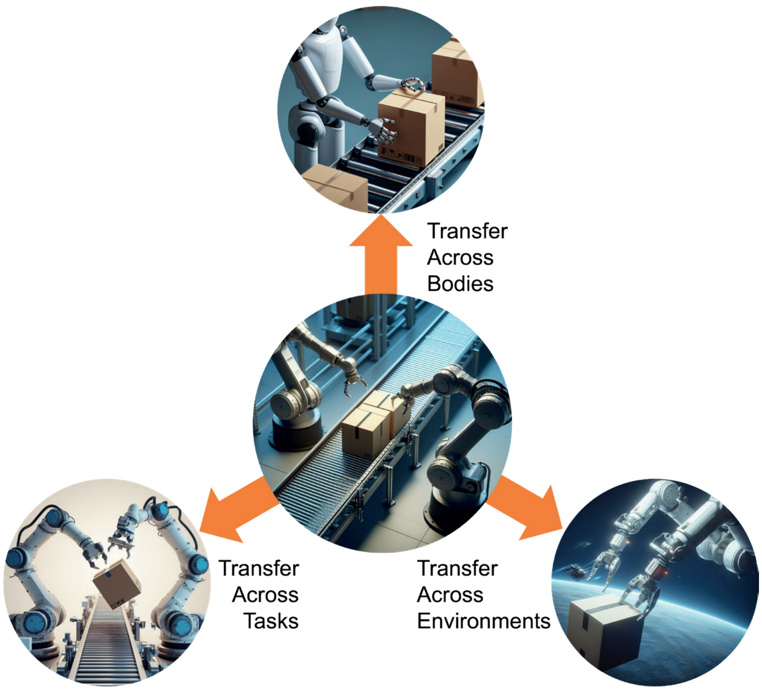


*Figure 2: The ability to transfer learning across robot bodies, tasks and environment is fundamental for achieving collaborations of different robots on one task). Schematic makes use of images generated by Microsoft Bing Creator.*

*Figure 3: Handing a package from a drone to a humanoid robot or single-arm robot manipulator requires to reconcile drastically different perception, from different viewpoints and sensors, and distinct robot actions from unimanual to bimanual actions (inspired by 2023 EuROBIN Hackaton). Image generated by Microsoft Bing Creator.*

**Transfer learning**: Future robots will need to be able to *transfer* what they learn: from one task to another, from one environment to another, and from one robot to another. Human intelligence relies on the ability to apply the knowledge acquired in one domain to new domains - thus solving new problems and facing unexpected situations – and to share knowledge among individuals. Similarly, robots need analytic and data-driven methods for learning skills from human demonstrations, transferring learned skills to novel tasks or different robots and environments, transferring skills learned in simulation to real robots, transferring learned perception routines between robots.

There are several open questions that need to be solved to reach transferrable robot learning. The first one is *what to transfer*: we need to develop criteria to select the learned knowledge about environments, objects, tasks constraints that can and should be transferred when dealing with new environments, objects, and tasks. A second question is *how to transfer*: for successful transfer to happen, prior knowledge of robot bodies may often be required, for example on sensors, kinematics, actuators, electronic hardware etc. Finally, we need to know *when to transfer*, developing algorithms to recognize similarities across environment, objects, tasks constraints, establishing if transfer of knowledge is at all possible in every specific case or if entirely new knowledge and novel learning cycles are needed.

## 6. CLOSING WORDS

Deployment of AI and robotics at large has become a more tangible target, possibly foreseeable in the next decade. AI has the potential to expand largely the capabilities and range of applications of robotics. At last, the multi-decade dream of intelligent, capable and useful robots is within sight. Major hurdles on the road include ensuring understandability and controllability for safe deployment and usage and achieving scalable, cost-effective solutions to support autonomy and resilience.

# REFERENCES


[1] A. Krizhevsky, I.Sutskever, and G. E. Hinton, 'ImageNet Classification with Deep Convolutional Neural Networks', Communications of the ACM 60, no. 6 (May 2017): 84–90. https://doi.org/10.1145/3065386

[2] B. Siciliano, O.Khatib (eds), Springer Handbook of Robotics, Springer, Cham (2016) https://doi.org/10.1007/978-3-319-32552-1

[3] B. Siciliano, L.Sciavicco, L.Villani, G. Oriolo, Robotics. Modelling, Planning and Control. Springer-Verlag London 2009. https://doi.org/10.1007/978-1-84628-642-1

[4] Roland Siegwart, Illah Reza Nourbakhsh and Davide Scaramuzza. Introduction to Autonomous Mobile Robots. The MIT Press, Cambridge, 2011. ISBN: 9780262015356.

[5] J. Gao, Z. Tao, N. Jaquier and T. Asfour, "K-VIL: Keypoints-Based Visual Imitation Learning," in IEEE Transactions on Robotics, vol. 39, no. 5, pp. 3888-3908, Oct. 2023, doi: .

[6] Harish Ravichandar et al., 'Recent Advances in Robot Learning from Demonstration', Annual Review of Control, Robotics, and Autonomous Systems 3, no. Volume 3, 2020 (3 May 2020): 297–330, https://doi.org/10.1146/annurev-control-100819-063206

[7] Learning Ambidextrous Robot Grasping Policies. Jeffrey Mahler, Mathew Matl, Vishal Satish, Mike Danielczuk, Bill DeRose, Stephen McKinley, Ken Goldberg. Science Robotics Journal. V4(26). Jan 2019

[8] Seungsu Kim, Ashwini Shukla, and Aude Billard, 'Catching Objects in Flight', IEEE Transactions on Robotics 30, no. 5 (October 2014): 1049–65, https://doi.org/10.1109/TRO.2014.2316022

[9] Antonio Loquercio et al., 'Learning High-Speed Flight in the Wild', Science Robotics 6, no. 59 (6 October 2021): eabg5810, https://doi.org/10.1126/scirobotics.abg5810

[10] Grollman, Daniel H., and Aude Billard. "Donut as i do: Learning from failed demonstrations." 2011 IEEE international conference on robotics and automation. IEEE, 2011.

[11] Chen, Letian, Rohan Paleja, and Matthew Gombolay. "Learning from suboptimal demonstration via self supervised reward regression." In Conference on robot learning, pp. 1262-1277. PMLR, 2021.

[12] Butepage, J., Black, M.J., Kragic, D. and Kjellstrom, H., 2017. Deep representation learning for human motion prediction and classification. In Proceedings of the IEEE conference on computer vision and pattern recognition (pp. 6158-6166).

[13] Calinon, Sylvain, and Aude Billard. "Active teaching in robot programming by demonstration." RO-MAN 2007-The 16th IEEE International Symposium on Robot and Human Interactive Communication. IEEE, 2007.

[14] Zollner, R., Tamim Asfour, and Rüdiger Dillmann. "Programming by demonstration: dual-arm manipulation tasks for humanoid robots." 2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)(IEEE Cat. No. 04CH37566). Vol. 1. IEEE, 2004.

[15] Nicolescu, Monica, and Maja J. Mataric. "Task learning through imitation and human-robot interaction." Models and mechanisms of imitation and social learning in robots, humans and animals: behavioural, social and communicative dimensions (2005): 407-424.

[16] Florence, Pete, et al. "Implicit behavioral cloning." Conference on Robot Learning. PMLR, 2022.

[17] Richard S. Sutton; Andrew G. Barto, Reinforcement Learning: An Introduction. MIT Press, Cambridge, 1998, ISBN: 9780262257053.

[18] Elia Kaufmann et al., 'Champion-Level Drone Racing Using Deep Reinforcement Learning', Nature 620, no. 7976 (August 2023): 982–87, https://doi.org/10.1038/s41586-023-06419-4

[19] Ilija Radosavovic et al., 'Real-World Humanoid Locomotion with Reinforcement Learning', Science Robotics 9, no. 89 (17 April 2024): eadi9579, https://doi.org/10.1126/scirobotics.adi9579

[20] Tuomas Haarnoja et al., 'Learning Agile Soccer Skills for a Bipedal Robot with Deep Reinforcement Learning', Science Robotics 9, no. 89 (10 April 2024): eadi8022, https://doi.org/10.1126/scirobotics.adi8022

[21] Ibarz, Julian, et al. "How to train your robot with deep reinforcement learning: lessons we have learned." The International Journal of Robotics Research 40.4-5 (2021): 698-721.

[22] Kober, Jens, and Jan Peters. "Imitation and reinforcement learning." IEEE Robotics & Automation Magazine 17.2 (2010): 55-62.

[23] Hester, Todd, et al. "Deep q-learning from demonstrations." Proceedings of the AAAI conference on artificial intelligence. Vol. 32. No. 1. 2018.

[24] Stephen Adams, Tyler Cody, and Peter A. Beling, 'A Survey of Inverse Reinforcement Learning', Artificial Intelligence Review 55, no. 6 (1 August 2022): 4307–46, https://doi.org/10.1007/s10462-021-10108-x

[25] Kim D, Kim S-H, Kim T, Kang BB, Lee M, Park W, et al. (2021) Review of machine learning methods in soft robotics. PLoS ONE 16(2): e0246102. https://doi.org/10.1371/journal.pone.0246102

[26] Subramanian Sundaram et al., 'Learning the Signatures of the Human Grasp Using a Scalable Tactile Glove', Nature 569, no. 7758 (May 2019): 698–702, https://doi.org/10.1038/s41586-019-1234-z

[27] Jeffrey Mahler and Ken Goldberg, 'Learning Deep Policies for Robot Bin Picking by Simulating Robust Grasping Sequences', in Proceedings of the 1st Annual Conference on Robot Learning (Conference on Robot Learning, PMLR, 2017), 515–24, https://proceedings.mlr.press/v78/mahler17a.html

[28] Open X.-Embodiment Collaboration et al., 'Open X-Embodiment: Robotic Learning Datasets and RT-X Models' (arXiv, 2 April 2024), https://doi.org/10.48550/arXiv.2310.08864

[29] Kaufmann, E., Bauersfeld, L., Loquercio, A. et al. Champion-level drone racing using deep reinforcement learning. Nature 620, 982–987 (2023). https://doi.org/10.1038/s41586-023-06419-4

[30] Ashish Kumar et al., 'RMA: Rapid Motor Adaptation for Legged Robots', in Robotics: Science and Systems XVII (Robotics: Science and Systems 2021, Robotics: Science and Systems Foundation, 2021), https://doi.org/10.15607/RSS.2021.XVII.011

[31] Jemin Hwangbo et al., 'Learning Agile and Dynamic Motor Skills for Legged Robots', Science Robotics 4, no. 26 (16 January 2019): eaau5872,

https://doi.org/10.1126/scirobotics.aau5872

[32] Ashish Vaswani et al., 'Attention Is All You Need', in Advances in Neural Information Processing Systems, vol. 30 (Curran Associates, Inc., 2017).

[33] Dhruv Shah et al., 'Navigation with Large Language Models: Semantic Guesswork as a Heuristic for Planning' (arXiv, 16 October 2023), https://doi.org/10.48550/arXiv.2310.10103

[34] Zhe Gan et al., 'Vision-Language Pre-Training: Basics, Recent Advances, and Future Trends' (arXiv, 17 October 2022), https://doi.org/10.48550/arXiv.2210.09263

[35] Brohan A. et al., 'RT-2: Vision-Language-Action Models Transfer Web Knowledge to Robotic Control' (arXiv, 28 July 2023), https://doi.org/10.48550/arXiv.2307.15818

[36] Toussaint, M., Logic-Geometric Programming: An Optimization-Based Approach to Combined Task and Motion Planning. Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI 2015).

[37] Semeraro, F. et al. "Human–robot collaboration and machine learning: A systematic review of recent research." Robotics and Computer-Integrated Manufacturing 79 (2023): 102432.

[38] Brunke, L, et al. "Safe learning in robotics: From learning-based control to safe reinforcement learning." Annual Review of Control, Robotics, and Autonomous Systems 5.1 (2022): 411-444.

[39] Hewing, L., et al. "Learning-based model predictive control: Toward safe learning in control." Annual Review of Control, Robotics, and Autonomous Systems 3.1 (2020): 269-296.

[40] Nghiem, Truong X., et al. "Physics-informed machine learning for modelling and control of dynamical systems." 2023 American Control Conference (ACC). IEEE, 2023.

[41] Tsukamoto, Hiroyasu, Soon-Jo Chung, and Jean-Jaques E. Slotine. "Contraction theory for nonlinear stability analysis and learning-based control: A tutorial overview." Annual Reviews in Control 52 (2021): 135-169.

[42] Dongho Kang et al., 'RL + Model-Based Control: Using On-Demand Optimal Control to Learn Versatile Legged Locomotion', IEEE Robotics and Automation Letters 8, no. 10 (October 2023): 6619–26, https://doi.org/10.1109/LRA.2023.3307008

[43] Jeffrey Ichnowski et al., 'Deep Learning Can Accelerate Grasp-Optimized Motion Planning', Science Robotics 5, no. 48 (18 November 2020): eabd7710, https://doi.org/10.1126/scirobotics.abd7710

[44] Thrun, Sebastian, and Tom M. Mitchell. "Lifelong robot learning." Robotics and autonomous systems 15.1-2 (1995): 25-46.

# 2. Control, Planning and Reasoning in the era of generative AI

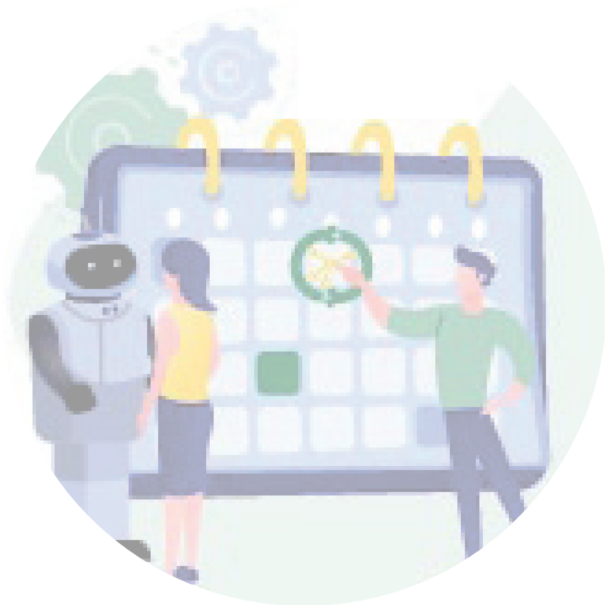

**Alin Albu-Schaeffer**, Institute of Robotics and Mechatronics, DLR-German Aerospace Center/Department of Informatics, Technical University of Munich, Germany

**Michael Beetz**, Institute for Artificial Intelligence, Computer Science Department, University of Bremen, Germany

**Lydia Kavraki,** Department of Computer Science, Rice University, Houston, USA

**Stefano Stramigioli**, University of Twente, Netherlands

**Bruno Siciliano**, Università di Napoli Federico II, Italy

**Cecilia Laschi**, National University of Singapore, Singapore

**Anibal Ollero**, Universidad de Sevilla, Spain

**Danica Kragic**, School of Computer Science and Communication, Royal Institute of Technology, Stockholm, Sweden

**Henrik I. Christensen**, UC San Diego, La Jolla, USA

**Aude Billard**, Learning Algorithms and Systems Laboratory (LASA), École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland

**Kazuhiro Kosuge**, University of Hong Kong

**Chiara Sabelli**, independent researcher

**Nicola Nosengo**, independent researcher

As it faces the overarching challenges of taking robots to complex and uncontrolled environments, executing complex and partially unpredictable tasks while operating autonomously and interacting more closely with humans, robotics will need to expand the current scope of planning, control, and reasoning techniques and combine them with generative AI and learning. This article looks at the main challenges in adapting control, planning, and reasoning to the next generations of robots. It defines a roadmap to move from the current state of the art to goals that can be reached in the short and medium term, and to open scientific challenges that will keep researchers occupied at least for the next ten to fifteen years. Short-term goals include the automation of new working environments, aerial manipulation, improved teleoperation, improved physics engines for simulation and semantic digital twins. Long-term challenges include model-based manipulation of deformable objects, model-based control of soft robots, new mathematical approaches to control and planning, long-horizon planning, multi-robot control, advanced reasoning for seamless collaboration with humans.

# 1. INTRODUCTION

Planning and controlling the movements of robots made of mechanical parts is a cornerstone problem – if not the cornerstone problem - of robotics. Industrial robotics was born when computer programs were first applied to guide the movement of robotic arms, mounted over a fixed base, in a tridimensional space. Over the years, the scope of control and planning techniques has expanded to include the movements of mobile robots that have no fixed base and can navigate an environment while performing tasks within it.

For both manipulators and mobile robots, the fundamental steps to be taken include:

a) creating a mathematical model of the robot itself, including its kinematic and dynamic behavior, and a mathematical model of the environment where it is situated[1]

b) finding a trajectory to move the robot from an initial configuration to a final, desired configuration, without colliding with the environment and respecting all the kinematic and dynamic constraints of the robot (motion planning)[2]

c) using actuators and sensors to create the motion required by the planned trajectory (control)

While the first generations of industrial robots were bound to follow predetermined trajectories, decades of development in motion planning and control have now led to robots that can adapt their trajectories in real-time, either to compensate for changes in the position of the object that the robot must handle, or to guarantee the safety of human operators in the vicinity of the robot[3]. In mobile robots, research and industrial development have led to control policies for wheeled, tracked, flying or swimming robot coupled with navigation algorithms that allow those robots to move autonomously in environments for which a precise map is available, and even to a certain degree in environments that they have never encountered before[4]. In certain cases, finding optimal trajectories according to user-specified criteria (e.g. length) is possible. Ultimately, to approximate the capabilities of living organisms and to operate without humans' supervision, robots need to be endowed with reasoning abilities that allow them to encode and use semantic knowledge to make inferences about the consequences of their actions, interpret situations never encountered before, and make decisions.

As it faces the overarching challenges of taking robots to complex and uncontrolled environments, executing complex and partially unpredictable tasks while operating autonomously and interacting more closely with humans, robotics will need to expand the current scope of planning, control and reasoning techniques. While data-driven techniques and machine learning are currently attracting much attention[5], there are several robotic applications that require predictable control and explainable behavior which can only be guaranteed using underlying models (of the robot, the task, the environment and of humans). And thanks to improvements in mathematical methods and computational technologies, there are still huge margins of advancement in model-based methods that do not rely primarily on learning.

The role of control, planning, and reasoning in shaping the future of intelligent robotics in the age of generative AI and foundation models, is closely tied to a broader philosophical question: do intelligent robots need to maintain explicit representations of their capabilities and bodies to operate effectively in human environments and accomplish human-scale tasks? Kahneman's dual-process theory of decision-making and intelligence offers a compelling framework for integrating generative AI with model-based techniques in a synergistic manner[6]. The theory posits two complementary modes of reasoning: System 1, which is fast, intuitive, and associative, and System 2, which is slower, deliberate, and analytical. In robotics, generative AI can embody the rapid, adaptive qualities of System 1, leveraging large-scale data and multimodal learning to predict actions and generate flexible behaviors in real time. Meanwhile, model-based techniques align with System 2, providing structured, logical reasoning grounded in explicit representations to ensure correctness, safety, and long-term planning.

The main goals of the next decade of research in these fields will be to exploit the current modelling techniques at the best of their potential, making planning and control faster and more robust; to improve modelling itself, and expand it to systems that have so far being considered intractable, such as soft robotic systems; to leverage new hardware and software tools for better and faster motion planning; to scale up planning and control methods to have multiple robots work collaboratively; to advance reasoning abilities in robot and combine them with learning and generative AI.

This article looks at the main challenges in adapting control, planning, and reasoning to the next generations of robots. It defines a roadmap to move from the current state of the art to low-hanging fruits that can be reached in the short and medium term, to open scientific challenges that will keep researchers occupied at least for the next ten to fifteen years of research (*see Figure 1*).

## 2. STATE OF THE ART

Six decades of work on industrial robotics and automation, both in academia and in the robotics industry, have led to well-established techniques for the interconnected problems of modelling (the kinematic analysis of the mechanical structure of a robot), motion planning (the generation of trajectories to take the robot from a given initial configuration to a desired final configuration), control (the realization of the desired motion by actuators and sensors ) and navigation (the ability of a mobile robot to know its position inside an environment and use it for planning and controlling new trajectories).

Model-based planning and control for manipulators operating in controlled environments and on rigid,

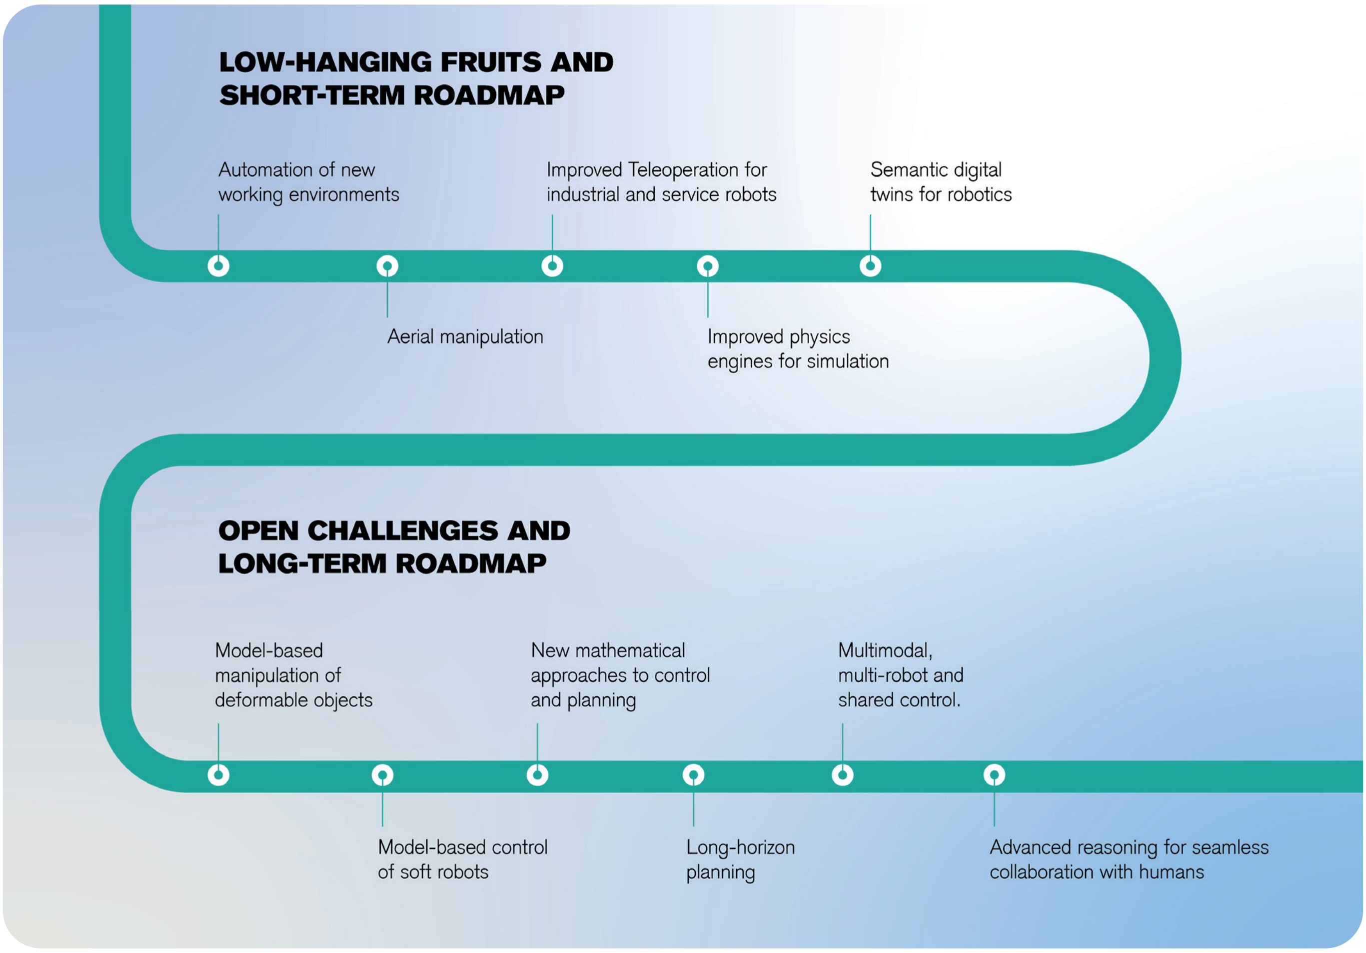


*Figure 1. Summary of short-term and long-term research goals presented in the article. The roadmap is not intended as a temporal sequence, but rather as a series of goals with increasing levels of complexity to be researched in parallel.*

a-priori known objects is solved and extensively deployed in the millions of robots operating in the automotive, chemical, electronic industry. Tridimensional, free-space motion planning with high degrees of freedom is fundamentally solved, allowing manipulators to work very efficiently in controlled environments such as assembly lines, warehouses, and automated laboratories[7]. A key advancement from the last two decades was the addition of on-line acquisition and elaboration of sensory feedback provided by visual, haptic, force/torque, laser sensors, that have allowed to develop more intelligent control algorithms for industrial robots based on principles such as proportional–integral–derivative (PID) control, adaptive control, tracking control[8 9].

Another essential contribution from research carried out in the recent decades has been compliant control, developed theoretically since the 1980s and deployed in the first two decades of the 21st century[10 11]. It allows controlling the energy injected in the systems and to allow interaction with humans rather than just following trajectories. Compliant control allows going from following trajectories to controlling the impedance and the interaction with humans and the environments. It has resulted in the development of co-bots (>link to HRI article) but has proven important also for classical industrial robots[12].

Hierarchical multitasking control has also greatly advanced at the research level. Control framework now exists that can create a priority scale of multiple tasks, making sure the tasks with lower priority are fulfilled only as long as they don't interfere with the highest priority task[13]. A humanoid robot could, for example, have a main task of manipulating an object, a secondary one of avoiding collisions with the environment and another of minimizing the effort. The methodical basis for these developments were laid down in the 1980s and 1990s and has been deployed during the last 20 years on humanoids or mobile robots equipped with arms.

Advancements have also been achieved in controlling complex dynamics that happen when the robots do not have a stable base attached to the ground, such as in the case of free-floating space robots equipped with manipulators, such as a satellite that can catch another satellite while floating in space and then stabilize the system. Algorithms have been introduced to control the center of mass, the momentum to or even free-floating robots with arms or in aerial manipulation, with an arm on-board a flying robot[14 15 16].

Teleoperation, e.g. the ability to control a robotic avatar over any distance in a closed loop, feeling the interaction forces and to achieve high-fidelity control, is state-of-the-art although not yet deployed on the large scale, and has been demonstrated in settings such as astronauts on the International Space Station controlling robots on the ground despite transmission delay and gravity effects[17] or deep sea operation of a humanoid underwater robot controlled by a human on the surface[18].

When it comes to mobile robots, 2-D path planning on flat terrains is largely a solved problem, including coverage path planning, that is the problem of computing the optimal path and project a collision-free trajectory to ensure the robot fully covers an area of interest within a certain time[19]. The problem of creating a map of a previously unknown environment, updating it continuously while keeping track of the robot's position in it has been solved by Simultaneous Localization and Mapping (SLAM) algorithm, that can operate efficiently in unknown cluttered environments using either lidars or cameras - or both[20 21]. Both coverage path planning and visual SLAM are successfully deployed in millions of robotic vacuum cleaners.

Kinodynamic trajectory optimization and control for wheeled in the absence of unexpected events is also state-of-the-art, as demonstrated by existing small-scale deployments of self-driving cars[22]. Multi Robot motion planning in known environments such as warehouses is achieved, at least when one is only concerned with the position of the robot and not with what is happening with its body[23]. Many aspects of the control of drones, especially quadcopters, have been solved and deployed. Autonomous navigation of flying robots using GNSS, or environment perception in partially denied environments where satellite signal is not available, is achieved. This allows controlling all the phases of the flight, including fully autonomous takeoff and landing even in constrained space, such as the perching of a flapping-wing robot on a branch[24]. Obstacle detection and avoidance is state of the art, both in indoor and benign outdoor environments, as long as there is no significant wind, which instead remains an unsolved problem for the control of drones.

Cognitive robot architectures exist that provide structured frameworks for integrating perception, reasoning, learning, and action capabilities within robotic systems, enabling them to exhibit intelligent, goal-oriented behavior. These architectures incorporate advanced cognitive reasoning mechanisms such as prospection, which allows robots to anticipate and simulate future scenarios; affordances, which enable the robot to perceive actionable possibilities within its environment; and attention, which helps prioritize relevant sensory information and tasks. They also include self-awareness for monitoring and adapting internal states, memory mechanisms like episodic and semantic memory for storing and leveraging past experiences, and situated reasoning for context-aware decision-making. Examples include the CRAM (Cognitive Robot Abstract Machine) architecture, designed to facilitate goal-directed tasks in everyday environments by using underdetermined plans that resolve into specific actions through runtime reasoning[25]; iSAC, that employs a multi-agent system combined with memory subsystems (e.g., sensory egosphere, semantic memory) and an internal rehearsal system for action simulation and selection[26]; and ARMAR which focuses on enabling humanoid robots to perform complex manipulation tasks in human environments by combining advanced perception, semantic reasoning, and action planning[27]. These architectures exemplify how cognitive frameworks enable robots to adapt to dynamic environments, interact with humans, and learn from experiences, bridging the gap between symbolic reasoning and sensorimotor interactions.

Knowledge representation and reasoning are fundamental to solve the robot body motion problem, enabling robots to infer and execute appropriate actions for complex tasks. This approach involves integrating symbolic representations, such as ontologies and axiomatic knowledge bases, to encapsulate information about objects, environments, and actions. For example, frameworks like KnowRob[28] and CRAM use such representations to infer the motion parameters required to achieve desired effects, such as grasping and lifting objects, while avoiding undesirable side effects like collisions or spillage. These systems combine abstract knowledge (e.g., social conventions or intuitive physics) with contextual information from sensors and episodic memories to adapt to dynamic and underdetermined scenarios, such as ”setting a table” in varying household environments. By leveraging structured reasoning mechanisms, robots can not only generate effective motion plans but also account for uncertainties and failures, thus enhancing reliability in open-ended domains.

## 3. RESEARCH PRIORITIES IN THE SHORT TERM AND LOW-HANGING FRUITS FOR INNOVATION

**Applying research results to new industrial domains.** Some important improvements in planning, control, and navigation would be relatively easy to achieve in the short term, as they are based on solutions that have been extensively studied in laboratories in the last decade and would mainly need additional effort (and adequate funding) to be scaled up, validated and commercially deployed.

A good example is the extension of automation to working environments that so far have remained only partially automated, such as large-scale research and testing laboratories. Laboratory automation with easy customization appears especially promising for chemistry labs, medical testing, and DNA sequencing. These contexts present varying degrees of complexity, ranging from tasks such as high throughput screening or quality control that are largely repetitive and can be addressed by traditional planning, control and programming, to more dynamic R&D setups in less structured environments that require novel robotics approaches with perception, knowledge representation, information sharing. Architecture models for the integration of existing solutions in life science laboratories in a plug-and-play fashion have been proposed and can be further developed[29]. The construction industry is another case where robots have a proven potential to improve productivity and enhance the safety of workers, and where new planning and control methods can be applied to tasks such as the on-site quality check and assembly of parts manufactured off-site.

For mobile robots, aerial manipulation has made great advancements and is a reality at the research level, with drones that can perform manipulation while flying or after perching to increase dexterity and force[30]; yet, deployment and commercial application require further technological development, in particular on expanding on-board real-time perception and planning capabilities to allow effective control of the forces exerted by – and felt by – the drone; and policy development, since the lack of a clear and consistent regulatory framework is currently constraining research and development.

Teleoperation for industrial and service robots is also ready for wider application thanks to the development of virtual reality and haptic feedback systems and can allow the automation of industrial processes (such as sanding, grinding, polishing) that are too complex for unmanned manipulators but where there is room to increase the safety and comfort of human workers, by physically removing them from the material being manipulated[31].

**Physics engines and digital twins as enabling technologies.** In terms of enabling technologies, hardware-accelerated motion planning for high-dimensional robots is close to the stage where it can be used for faster predictive control. Current motion planners can solve realistic and challenging problems in hundreds of milliseconds to dozens of seconds on consumer CPUs, which is too slow for reactive operation in evolving environments and prevents the achievement of higher-level autonomy. However, several strategies to accelerate motion planning have been demonstrated, typically combining some parallelisation of computation with hardware acceleration. While GPU-based acceleration implies a huge computational cost and introduces latency in communication, efficient acceleration can also be achieved on ordinary CPUs by exploiting some of their native features. Motion planning time for manipulators could thus be reduced to microseconds, significantly accelerating the movements of industrial arms with more than 7 DoFs[32]. Advancement in on-the-fly motion planning would also facilitate the development of socially adequate robots – not able to fully cooperate with humans but that can have limited interaction in structured environments like hospitals.

A key area of effort in the short and medium term must be the improvement of robotic simulators. These are key tools for modelling, motion planning, and control, as they allow testing planning and control algorithms safely and inexpensively before trying them out in the real world and on a real robot. The so-called sim-to-real gap is currently a limiting factor not only for data-driven and learning-based approaches, but also for model-based planning and control. Increasing the fidelity of physics simulators and physics engines is crucial to overcome this gap. Thanks in part to huge investments from the gaming industry, better physics engines and visual rendering of physical interaction are now becoming available and can be transferred to the robotic domain with relatively easy adaptation to obtain robotics-enabling simulations that are computationally lighter, modular, faster, and more resource-efficient than current ones[33]. However, for them to be used in robotics research it is important that different physics engines can work together by relying on unified modelling abstraction and hierarchies. Additionally, the needs and the objectives of physics simulation for robotics are very different from those of virtual reality and computer gaming, and more work is needed to define the ideal trade-offs between fidelity of the simulation, computational cost, and usefulness in helping define tractable control policies.

Digital twins go beyond pure simulation and modelling by creating a bidirectional interaction between the virtual and the physical. A digital twin can be defined as a “a set of virtual information constructs that mimics the structure, context, and behavior of a natural, engineered, or social system (or system-of-systems), is dynamically updated with data from its physical twin, has a predictive capability, and informs decisions that realize value"[34]. Digital twins of large natural and man-made environments, including factories, can be developed to enable robot programming at a more abstract level and facilitate the realization of robotics tasks of much larger complexity.

Semantic digital twins will substantially advance robotic capabilities by embedding rich, structured knowledge into digital representations of physical environments. These twins integrate detailed 3D models with semantic annotations, enabling robots to access context-specific information about objects, their properties, and relationships. For example, a robot can query a semantic digital twin to determine the precise 6D pose of a handle or understand the articulation model of a cupboard door, allowing it to create motion planning and control problems automatically. By serving as virtual knowledge bases, semantic digital twins not only provide data but also compute truth values of relationships dynamically, transforming how robots plan and execute tasks.

Ongoing research in generative AI is addressing the automated construction of semantic digital twins. This advancement has the potential to bridge the gap between data-driven and model-based approaches, enhancing robots' ability to create accurate digital representations of environments autonomously. As this technology matures, it will further enrich semantic digital twins, enabling robots to operate with greater precision, adaptability, and context-awareness across diverse applications.

# 4. OPEN CHALLENGES FOR A LONG-TERM ROADMAP

**Model-based manipulation of deformable objects.** An open challenge that robotics will need to address in the next decade, and most likely extending well beyond, is how to advance the model-based motion planning and control of soft systems. This problem covers two interrelated but distinct challenges.

The first one is the manipulation of deformable objects, that is currently a challenge for industrial robots and yet would be crucial for many applications, from medicine to agriculture to the automation of the textile and food sectors[35]. Progress will be needed on the hardware side, with the design of new soft grippers but also on the modelling, planning and control side. While machine learning has a significant potential in this regard, it is unlikely that industrial deployment of manipulation of deformable objects can be based on ML alone, especially in safety-critical applications that require predictability and explainability of robot performance.

Several methods exist in the literature for modelling the behavior of deformable objects[36]. Examples include mass-spring systems, position-based dynamics, and continuum mechanics. All have been applied with varying success at the experimental to cases including food, tissues, fabric, paper, and each has its own limitations. These methods have also been used to create physics-based simulators such as SOFA, PhysX, MuJoCo, that provide development environments for state estimation and motion planning in manipulation tasks involving deformable object, and that in turn allow planning and control approaches for deformable objects. Future research will need to evolve these methods and define the best mix of techniques to tackle specific manipulation problems – be it folding a shirt, in-hand manipulation of sponge-like objects, or tying a rope.

**Model-based control of soft robots.** The other challenge is the development of analytical models of robots that are themselves soft. The development of control algorithms in soft robotics has followed a reversed path compared to most domains in computation and robotics. For soft robots, learning-based approaches have been applied[37 38] before model-based ones, which were viewed as too challenging, or simply not applicable because of the virtually infinite degrees of freedom of a soft robotic system interacting with unpredictable environments. More recently though, researchers have found more and more effective ways to approximate soft robotics dynamics, paving the way for new modelling approaches[39]. In some cases, even simplified models have been shown to improve the performance with respect to model-free approaches[38 40]. For example, many soft robots have one dimension longer than the other two, and their whole configuration can be simulated by only considering deformations along that axis, vastly simplifying the problem. When that is not possible, new mathematical approaches – such as finite-dimensional modelling techniques that use partial differential equations to simulate infinite-dimensional systems – have been developed, that combine computational tractability with enough precision to describe the behavior of soft robots. By translating the volume of the robot into a mesh, that is a set of nodes and the information on their neighbors, it becomes possible to approximate the entire volume of the soft system using interpolation. The choice of the modelling technique determines the best control strategy to be used for the robot, which may focus on curvature, strain, or volume control, and that may or may not combine actuation and under-actuation. Research in the next decade needs to focus on further developing and testing modelling techniques and model-based planning/control for soft robots, including soft aerial robots, as well as understanding how they can be optimally combined with data-driven and learning-based techniques[41].

**New mathematical approaches to control and planning.** For both soft and rigid robots, a promising avenue of research is to look beyond the traditional approach to robot motion generation - that is to first plan trajectories on a kinematic level and then develop controllers for tracking the planned trajectories - taking the robot hardware as a priori. The study of intrinsic robot dynamics can translate into methodologies to generate highly efficient motions. New solutions are available to extend the methodical basis for modeling and controlling them, such as geometric mechanics and dynamics, differential geometry, and algebraic topology, that can mathematically describe the nonlinear oscillations that a robotic system may have. There are many examples of robotic motions, such as galloping or bouncing, that could be realized by exploiting intrinsic oscillation modes rather than being enforced on the system, in other words designing the robot so that it favours the desired set of movements[42].

For the next generation of mobile robots, fundamental research at the intersection with physics will be needed on how to effectively model the interaction

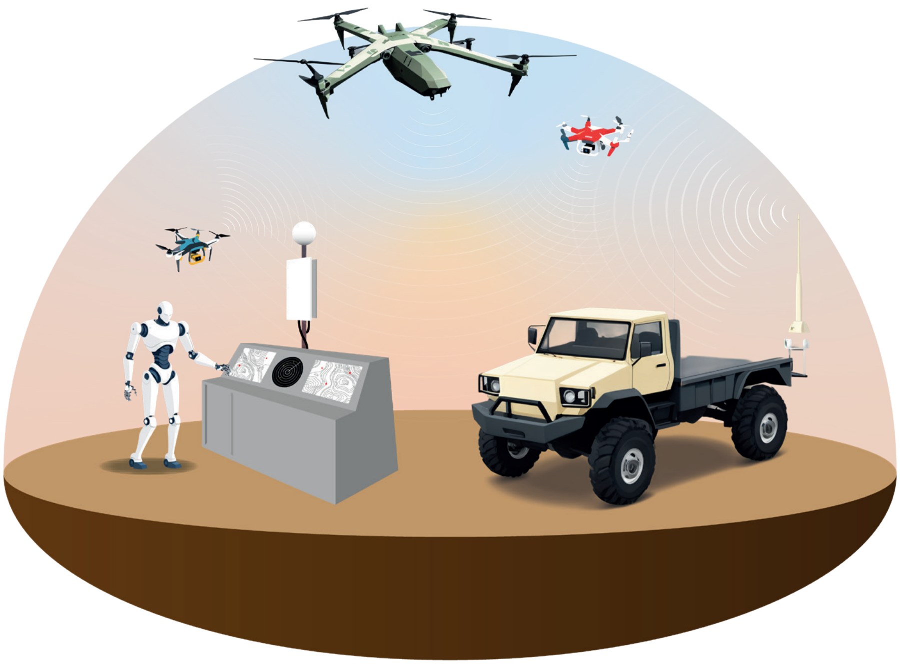

*Figure 2. An open challenge is how to control multi-robot systems where several robots of different types co-operate on tasks and share representations of the environment. that they may observe from varying points of view and with different sensors.*

of robots with their environment, and especially the complex case of interaction with fluids such as air, water, viscous substances). For example, in flying robotics, complex aerodynamic modelling will be needed to predict the unsteady lift and trust generated by a fixed-winged robotic bird because of the interaction with vortex formation around the wing, and particularly a flapping wing[43, 44]. Similar cases can be made for marine robotics, or for robots that have to move in viscous fluids such as oil or mud, or dig into sand and soil, to act autonomously in complex and extreme environments without human supervision.

**Long-horizon planning.** Long-horizon planning – the ability to consider action consequences over a long temporal period when moving towards a symbolically specified goal, a mission rather than merely a target position - is a necessary requisite for autonomous behavior in robots, but as of today it is still an open challenge because of computational cost and of the intrinsic difficulty in planning beyond a few short-term steps in realistic application settings[45]. Work is ongoing on new theoretical approaches to long-horizon planning – such as incorporating abstract strategies in task-planning routines and evaluating their affordance – that allow to practically accelerate long-horizon planning, with the goal of making it a tractable problem in realistic use cases.

Another significant open challenge will be fast motion planning under uncertainty, that requires computational approaches that incorporate from the beginning the uncertainty of the environment in motion planning algorithms[46]. Here the Finite Element Method (FEM) is proving useful in generating high-quality motion plans for use cases involving deformable objects, such as guiding steerable needles through deformable tissue for minimally invasive biopsies and drug-delivery, and manipulating planar tissues to align interior points at desired coordinates for precision treatment.

*Figure 3. A cognitive robot tasked with preparing a bowl of cereal for breakfast would face challenges with implicit knowledge that allows humans to do the same task without explicit planning. It would need to know that "a bowl of cereal" implies the use of milk, or what container can be used as the bowl or where to find the cereal. Without common-sense knowledge that provides answers to these challenges, the robot may search the whole kitchen for milk instead of starting with the most probable location (the fridge) or it would not understand that a found container could be used as the bowl (adapted from reference 49).*

**Multimodal, multi robot and shared control.** For legged robots a key open challenge is planning and control of multimodal locomotion, that allows the robot to switch between walking, climbing, jumping, and squeezing through narrow passages. Fundamental interdisciplinary research will be needed to understand and model how living organisms achieve effective multimodal locomotion, and this bioinspiration will be key to understand how to integrate open-loop and closed-loop control, active and passive control, model-based and learning-based strategies to achieve multimodal mobile robots[47].

Significant improvements are needed on control of multi-robot systems where several robots of different types co-operate on tasks, including control of robotics swarms with several tens, or even hundreds, of individuals: here, progress will be required on creating shared representations of the environment that different robots may observe from varying points of view and with different sensors (see Figure 2) as well as on common operative systems and communication protocol that go beyond current standards[48].

Shared control between humans and autonomous agents will be another important area of research. Many future application scenarios will require robots combining an autonomous agent, which controls part of the robot and a human controlling the rest. For example, think of an assistive device made of a wheelchair equipped with a robotic arm. The wheelchair would have 3 degrees of freedom and the arm may have 8 additional ones, but the human should not be in charge of controlling all 11 DoFs. She or he may be required to act on the main ones, indicating a direction or a desired action, and then the control system would need to take over and stabilize all other degrees of freedom.

**Semantic reasoning for robots.** Autonomous behaviour requires robots to have reasoning abilities to interpret their environment and cope with new and underdetermined tasks, new environments or new objects. Achieving this goal implies equipping robots with commonsense knowledge including physics, causality, objects with their locations, properties and relationships, the psychology of human beings - a so-called computational Theory of Mind[49].

The ultimate challenge is robots jointly accomplishing tasks for, and together with humans, as it is needed for robots that are co-workers of humans and robots that empower people to improve their quality of life. For this, robots require advanced reasoning capabilities that enable seamless collaboration in shared tasks. These capabilities include negotiating roles in joint activities, dynamically allocating subtasks, and adapting to human feedback to ensure alignment with shared goals. Robots must maintain a robust understanding of the context of the task, which involves reasoning about human intentions, the current state of the environment, and their own operational constraints. By integrating task planning with real-time feedback, robots can effectively co-construct actions with humans, ensuring mutual understanding and efficiency[50].

In addition to negotiation, robots must be capable of both giving and receiving help during task execution. This involves reasoning about when humans might require assistance, proactively offering help, and coordinating their actions without disrupting human efforts. Equally important is the ability to ask for help when needed, which requires self-awareness of their own limitations and the ability to articulate specific needs clearly. Robots must dynamically switch between autonomous operation and guided intervention, leveraging human input to overcome gaps in knowledge or capability. These interactions rely on the robot's ability to simulate potential actions, predict outcomes, and adjust its behavior to ensure that joint tasks proceed smoothly and efficiently.

Underlying these reasoning capabilities is the need for robust knowledge representation and decision-making frameworks (Figure 3). Robots must represent objects, events, and[51] relationships in structured formats that support real-time reasoning, enabling them to model their capabilities and limitations accurately. Incorporating probabilistic reasoning allows robots to operate under uncertainty, adapt to changes in the environment, and learn from both successes and failures. By combining intuitive physics and commonsense reasoning with task-specific knowledge, robots can anticipate human needs, avoid undesired outcomes, and continuously improve their performance in joint activities.

## 5. CLOSING WORDS

Control, planning, and reasoning have provided the foundations of robotics, and will remain central also in the age of deep learning and generative AI, shaping the future of intelligent robots. Intelligence ultimately involves maintaining representations and reasoning about them, and explicit models enable rigorous computational frameworks. The significance of such approaches lies in their ability to bypass data dependency and provide results that are correct, transferable, generalizable, and optimal. This ensures safety, trustworthiness, and reliability—qualities indispensable for robots operating in dynamic, human-centric environments.

# REFERENCES

[1] J. Hollerbach, W. Khalil, and M. Gautier, 'Model Identification', in Springer Handbook of Robotics, ed. Bruno Siciliano and Oussama Khatib (Berlin, Heidelberg: Springer, 2008), 321–44, https://doi.org/10.1007/978-3-540-30301-5_15.

[2] H. Choset, K.M. Lynch, S. Hutchinson, G. Kantor, W. Burgard, L.E. Kavraki, S. Thrun: Principles of Robot Motion: Theory, Algorithms, and Implementations (MIT Press, Cambridge 2005).

[3] A. M. Annaswamy and A. L. Fradkov, 'A Historical Perspective of Adaptive Control and Learning', Annual Reviews in Control 52 (1 January 2021): 18–41, https://doi.org/10.1016/j.arcontrol.2021.10.014.

[4] A. Loganathan and N.Syazreen Ahmad, 'A Systematic Review on Recent Advances in Autonomous Mobile Robot Navigation', Engineering Science and Technology, an International Journal 40 (1 April 2023): 101343, https://doi.org/10.1016/j.jestch.2023.101343.

[5] Q. Vuong et al. "Open x-embodiment: Robotic learning datasets and RT-x models." Towards Generalist Robots: Learning Paradigms for Scalable Skill Acquisition@ CoRL2023. 2023.

[6] D. Kahneman, Thinking, Fast and Slow. New York : Farrar, Straus and Giroux, 2011.

[7] A. Orthey, C. Chamzas, and L. E. Kavraki, 'Sampling-Based Motion Planning: A Comparative Review', Annu. Rev. Control Robot. Auton. Syst., vol. 7, no. Volume 7, 2024, pp. 285–310, Jul. 2024, doi: 10.1146/annurev-control-061623-094742.

[8] J. Craig, Ping Hsu, and S. Sastry, 'Adaptive Control of Mechanical Manipulators', in 1986 IEEE International Conference on Robotics and Automation Proceedings, vol. 3, 1986, 190–95, https://doi.org/10.1109/ROBOT.1986.1087661.

[9] Y. Choi and W.K. Chung, eds., PID Trajectory Tracking Control for Mechanical Systems, vol. 298, Lecture Notes in Control and Information Sciences (Berlin, Heidelberg: Springer, 2004), https://doi.org/10.1007/978-3-540-40041-7.

[10] C. Ott et al., 'A Passivity Based Cartesian Impedance Controller for Flexible Joint Robots - Part I: Torque Feedback and Gravity Compensation', in IEEE International Conference on Robotics and Automation, 2004. Proceedings. ICRA '04. 2004, vol. 3, 2004, 2659-2665 Vol.3, https://doi.org/10.1109/ROBOT.2004.1307462.

[11] F. Ficuciello, L. Villani, and B. Siciliano, 'Variable Impedance Control of Redundant Manipulators for Intuitive Human–Robot Physical Interaction', IEEE Transactions on Robotics 31, no. 4 (August 2015): 850–63, https://doi.org/10.1109/TRO.2015.2430053.

[12] T. Zhang et al., 'A Review of Compliant Control for Collaborative Robots', in 2021 IEEE 16th Conference on Industrial Electronics and Applications (ICIEA), 2021, 1103–8, https://doi.org/10.1109/ICIEA51954.2021.9516193.

[13] X. Wu et al., 'Passive Decoupled Multitask Controller for Redundant Robots', IEEE Transactions on Control Systems Technology 31, no. 1 (January 2023): 1–16, https://doi.org/10.1109/TCST.2022.3162990.

[14] F. Basana et al., 'Satellite and Robotic Arm Combined Control for Spacecraft Close-Proximity Operations', CEAS Space Journal, 3 August 2024, https://doi.org/10.1007/s12567-024-00560-0.

[15] A. M. Giordano, C. Ott, and A. Albu-Schäffer, 'Coordinated Control of Spacecraft's Attitude and End-Effector for Space Robots', IEEE Robotics and Automation Letters 4, no. 2 (April 2019): 2108–15, https://doi.org/10.1109/LRA.2019.2899433.

[16] F. Ruggiero, V. Lippiello and A. Ollero, "Aerial Manipulation: A Literature Review," in IEEE Robotics and Automation Letters, vol. 3, no. 3, pp. 1957-1964, July 2018, doi: 10.1109/LRA.2018.2808541.

[17] B. Weber et al., 'Teleoperating Robots from the International Space Station: Microgravity Effects on Performance with Force Feedback', in 2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2019, 8144–50, https://doi.org/10.1109/IROS40897.2019.8968030.

[18] G. Brantner and O. Khatib, 'Controlling Ocean One', in Field and Service Robotics, ed. Marco Hutter and Roland Siegwart (Cham: Springer International Publishing, 2018), 3–17, https://doi.org/10.1007/978-3-319-67361-5_1.

[19] C. S. Tan, R. Mohd-Mokhtar, and M. R. Arshad, 'A Comprehensive Review of Coverage Path Planning in Robotics Using Classical and Heuristic Algorithms', IEEE Access 9 (2021): 119310–42, https://doi.org/10.1109/ACCESS.2021.3108177.

[20] C. Stachniss, J. J. Leonard, and S. Thrun, 'Simultaneous Localization and Mapping', in Springer Handbook of Robotics, ed. Bruno Siciliano and Oussama Khatib, Springer Handbooks (Cham: Springer International Publishing, 2016), 1153–76, https://doi.org/10.1007/978-3-319-32552-1_46.

[21] C. Cadena et al., 'Past, Present, and Future of Simultaneous Localization and Mapping: Toward the Robust-Perception Age', IEEE Transactions on Robotics 32, no. 6 (December 2016): 1309–32, https://doi.org/10.1109/TRO.2016.2624754.

[22] C. Rösmann, F. Hoffmann, and T. Bertram, 'Kinodynamic Trajectory Optimization and Control for Car-like Robots', in 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2017, 5681–86, https://doi.org/10.1109/IROS.2017.8206458.

[23] C. Mangette and P. Tokekar, 'Multi-Robot Coordinated Planning in Confined Environments under Kinematic Constraints', in 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2020, 7999–8004, https://doi.org/10.1109/IROS45743.2020.9341213.

[24] R. Zufferey, J. Tormo-Barbero, D. Feliu-Talegón, S.R. Nekoo, J.A. Acosta, and A. Ollero. "How ornithopters can perch autonomously on a branch". Nature Commun 13, 7713 (2022).

[25] M. Beetz, L. Mösenlechner and M. Tenorth, "CRAM — A Cognitive Robot Abstract Machine for everyday manipulation in human environments," 2010 IEEE/RSJ International Conference on Intelligent Robots and Systems, Taipei, Taiwan, 2010, pp. 1012-1017, doi: 10.1109/IROS.2010.5650146.

[26] Y. Xiao et al., "Reasoning Over the Air: A Reasoning- Based Implicit Semantic-Aware Communication Framework," in IEEE Transactions on Wireless Communications, vol. 23, no. 4, pp. 3839-3855, April 2024, doi: 10.1109/TWC.2023.3312115

[27] K. Berns, T. Asfour, R. Dillmann, (2000). Design and Control of the Humanoid Robot ARMAR. In: Morecki, A., Bianchi, G., Rzymkowski, C. (eds) Romansy 13. International Centre for Mechanical Sciences, vol 422. Springer, Vienna. https://doi.org/10.1007/978-3-7091-2498-7_32

[28] M. Tenorth and M. Beetz, "KNOWROB — knowledge processing for autonomous personal robots," 2009 IEEE/RSJ International Conference on Intelligent Robots and Systems, St. Louis, MO, USA, 2009, pp. 4261-4266, doi: 10.1109/IROS.2009.5354602.

[29] Á. Wolf et al., 'Towards Robotic Laboratory Automation Plug & Play: Reference Architecture Model for Robot Integration', SLAS Technology 29, no. 4 (1 August 2024): 100168, https://doi.org/10.1016/j.slast.2024.100168.

[30] A. Ollero, M. Tognon, A. Suarez, D. Lee and A. Franchi, "Past, Present, and Future of Aerial Robotic Manipulators," in IEEE Transactions on Robotics, vol. 38, no. 1, pp. 626-645, Feb. 2022, doi: 10.1109/TRO.2021.3084395

[31] C. González et al., 'Advanced Teleoperation and Control System for Industrial Robots Based on Augmented Virtuality and Haptic Feedback', Journal of Manufacturing Systems 59 (1 April 2021): 283–98, https://doi.org/10.1016/j.jmsy.2021.02.013.

[32] W. Thomason, Z. Kingston, and L. E. Kavraki, 'Motions in Microseconds via Vectorized Sampling-Based Planning', in 2024 IEEE International Conference on Robotics and Automation (ICRA), 2024, 8749–56, https://doi.org/10.1109/ICRA57147.2024.10611190.

[33] 'A Review of Physics Simulators for Robotic Applications | IEEE Journals & Magazine | IEEE Xplore', accessed 3 November 2024, https://ieeexplore.ieee.org/document/9386154.

[34] Foundational Research Gaps and Future Directions for Digital Twins. Washington, D.C.: National Academies Press, 2024. doi: 10.17226/26894.

[35] V. E. Arriola-Rios et al., 'Modeling of Deformable Objects for Robotic Manipulation: A Tutorial and Review', Frontiers in Robotics and AI 7 (17 September 2020), https://doi.org/10.3389/frobt.2020.00082.

[36] H. Yin, A. Varava, and D. Kragic, 'Modeling, Learning, Perception, and Control Methods for Deformable Object Manipulation', Science Robotics 6, no. 54 (26 May 2021): eabd8803, https://doi.org/10.1126/scirobotics.abd8803.

[37] T. G. Thuruthel, Y. Ansari, E. Falotico, and C. Laschi. "Control Strategies for Soft Robotic Manipulators: A Survey", Soft Robotics 5, no. 2 (1 April 2018): 149 - 163, https://doi.org/10.1089/soro.2017.0007.

[38] C. Laschi, T. G. Thuruthel, F. Lida, R. Merzouki and E. Falotico, "Learning-Based Control Strategies for Soft Robots: Theory, Achievements, and Future Challenges," in IEEE Control Systems Magazine, vol. 43, no. 3, pp. 100-113, June 2023, doi: 10.1109/MCS.2023.3253421.

[39] G. Mengaldo, F. Renda, S. L. Brunton, M. Bächer, M. Calisti, C. Duriez, G. S. Chirikjian, and C. Laschi, "A concise guide to modelling the physics of embodied intelligence in soft robotics", Nature Reviews Physics 4, no. 9 (2022): 595-610.

[40] C. Della Santina, C. Duriez, and D. Rus, 'Model-Based Control of Soft Robots: A Survey of the State of the Art and Open Challenges', IEEE Control Syst. Mag., vol. 43, no. 3, pp. 30–65, Jun. 2023, doi: 10.1109/MCS.2023.3253419.

[41] C. A. Monje Micharet and C. Laschi, 'Editorial: Advances in Modeling and Control of Soft Robots', Frontiers in Robotics and AI 8 (21 May 2021), https://doi.org/10.3389/frobt.2021.706514.

[42] A. Albu-Schäffer and A. Sachtler, 'What Can Algebraic Topology and Differential Geometry Teach Us About Intrinsic Dynamics and Global Behavior of Robots?', in Robotics Research, ed. A. Billard, T. Asfour, and O. Khatib (Cham: Springer Nature Switzerland, 2023), 468–84, https://doi.org/10.1007/978-3-031-25555-7_32.

[43] C. Ruiz, J.Á. Acosta, A. Ollero. "Aerodynamic reduced-order Volterra model of an ornithopter under high-amplitude flapping". Aerospace Science and Technology. Vol 121, 107331 , Elsevier, 2022

[44] L. G. Groot Koerkamp et al., 'Leading Edge Vortex Formation on a Flapping Wing Robotic Bird', in AIAA AVIATION 2023 Forum, AIAA AVIATION Forum (American Institute of Aeronautics and Astronautics, 2023), https://doi.org/10.2514/6.2023-4243.

[45] K. Elimelech et al., 'Accelerating Long-Horizon Planning with Affordance-Directed Dynamic Grounding of Abstract Strategies', in 2024 IEEE International Conference on Robotics and Automation (ICRA), 2024, 12688–95, https://doi.org/10.1109/ICRA57147.2024.10610486.

[46] H. Durrant-Whyte, N. Roy, and P. Abbeel, 'Motion Planning under Uncertainty in Highly Deformable Environments', in Robotics: Science and Systems VII (Robotics: Science and Systems VII, MIT Press, 2012), 241–48, https://ieeexplore.ieee.org/document/6301045.

[47] A. J. Ijspeert, 'Biorobotics: Using Robots to Emulate and Investigate Agile Locomotion', Science 346, no. 6206 (10 October 2014): 196–203, https://doi.org/10.1126/science.1254486.

[48] M. Dorigo et al., Reflections on the future of swarm robotics. Sci. Robot. 5, eabe4385(2020).10.1126/scirobotics.abe4385.

[49] J.-P. Töberg et al., 'Commonsense Knowledge in Cognitive Robotics: A Systematic Literature Review', Frontiers in Robotics and AI 11 (29 February 2024), https://doi.org/10.3389/frobt.2024.1328934.

[50] M. Beetz, R. Chatila, J. Hertzberg, & F. Pecora, (2016). AI reasoning methods for robotics. Springer Handbook of Robotics, 329-356.

# 3. Human-Robot Interaction: Successes, Hurdles and Remaining Challenges

**Aude Billard**, Learning Algorithms and Systems Laboratory (LASA), École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland

**Astrid Rosenthal-von der Pütten**, Aachen University, Germany

**Serena Ivaldi**, Inria Center of the University of Lorraine, France

**Alin Albu-Schaeffer**, Institute of Robotics and Mechatronics, DLR-German Aerospace Center/Department of Informatics, Technical University of Munich, Germany

**Rachid Alami**, Laboratory for Analysis and Architecture of Systems (LAAS-CNRS) Toulouse, France

**Tamim Asfour**, Karlsruhe Institute of Technology, Germany

**Gordon Chen**g, Technical University of Munich, Germany

**Christophe Leroux**, CEA, France

**Danica Kragic**, School of Computer Science and Communication, Royal Institute of Technology, Stockholm, Sweden

**Nicola Nosengo**, independent researcher

**Chiara Sabelli**, independent researcher

The past decades have seen an increasing number of robots deployed in the vicinity of humans, from vacuum cleaners roaming in our living rooms, drones flying over our heads, to prostheses attached to our bodies. To increase trust and reduce risks, it is urgent and necessary that robots become cognizant of their environment and *socially aware*. They must be able to interpret, predict and reason about both human behavior and their own behavior.

This article aims to summarize existing solutions and open challenges over the next two decades towards the development of robotic applications capable of interacting with humans in a pertinent and helpful manner in any environment. Such applications can help tackle societal challenges, from assisting an aging population to monitoring the environment in order to mitigate and adapt to the effects of climate change and managing the impacts of natural hazards, such as earthquakes and floods. The article reviews successful examples of robots interacting and supporting humans, and delineates which breakthroughs, both in modeling and technology, have allowed such applications. It then highlights the low-hanging fruits, technologies that could improve the quality, effectiveness and versatility of the interaction and collaboration between humans and robots in the short- and medium- term that do not require scientific breakthroughs but rather clever strategies of technology transfer. It ends with a discussion of long-term scientific challenges that will require novel and interdisciplinary efforts to fulfill the vision of human-centered robotics.

# 1. INTRODUCTION

Today, all efforts globally are turned toward designing the next generation of robots, that of robots that will be employed and function in close or direct interactions with lay users. We are no longer in the realm of factory robots used by well-trained practitioners. We seek the design of autonomous wheelchairs, smarter and more dexterous prostheses and new drones and personal mobility devices that can navigate autonomously and safely to our doorsteps and on pedestrian lanes. It is not conceivable that these robots be programmed without a deep understanding of the social, ethical, cultural rules that underpin human environments.

Developing robots that are cognizant of the world that surrounds them has led to a wide range of efforts worldwide, all of which fall under the general field of human-robot interaction (HRI). The scope of HRI spans from developing algorithms and interfaces to facilitate seamless interaction between humans and robots, to conducting observations and experimental evaluations of how stakeholders utilize robots in various contexts. It encompasses both non-physical interaction—such as verbal and gesture-based communication—and physical interaction—where robots and humans are either in contact, as in prostheses, or in indirect contact, as when they jointly carry an object.

Originally an offspring of Human-Computer Interaction, HRI became a research field of its own in the mid-1990s, gaining increasingly more attention over the following three decades. It established itself with the launch of the IEEE-ACM International Conference on HRI in 2005 and subsequently of a few dedicated journals. HRI is fundamentally an interdisciplinary field of research, and requires close collaboration between roboticists and social scientists, cognitive scientists, psychologists, economists and philosophers. Their expertise is crucial to model human behavior and develop robots capable of interpreting and predicting the actions of the humans they interact with. It is also crucial to make sure that robots behave and speak in ways that are socially adequate and effective when communicating and collaborating with humans.

# 2. INTERACTION TYPES, SENSORS, AND INTERFACES FOR HRI

Human-robot interaction can be either haptic/physical–when humans and robots get into actual contact with each other–or non-physical. Physical human-robot interaction is being used extensively for teleoperation or for teaching robots to perform tasks through what is known as programming by demonstration or learning from demonstration. Moreover, modern collaborative robots (cobots) are designed to work together with humans, for example as a “third hand”[1] or for jointly manipulating large and heavy objects[2].

Non-physical interaction can be both verbal, when humans instruct robots what to do, and non-verbal, for instance by using robot’s eye gaze[3] to augment verbal communication, convey the robot’s internal state, or gather the attention of the human, possibly directing it toward an object involved in a joint task. This builds on the unique human ability, surpassing that of non-human primates, to infer others’ intentions from eye gaze. Physical and non-physical HRI can be combined in the most complex systems and tasks[4].

Over the last three decades, the type and complexity of both physical and non-physical HRI have evolved significantly due to several factors. First the introduction of new materials and new sensors, such as artificial skins and haptic interfaces, has enhanced robot’s ability to detect and interpret physical contacts with humans. Second, the design of more realistic human-like bodies, such as androids and human-like avatars, has enabled HRI to mimic certain aspects of human-human interaction. Finally, recent advancements in speech recognition and Large Language Models (LLMs) have greatly improved verbal interactions with robots, enabling more complex dialogues.

**Physical HRI**. In physical human-robot interaction, the key achievement that allowed the transition from the classical rigid, fixed-base robots to those capable of safely interacting with humans is compliant control[5], that is the possibility to regulate the energy injected in the systems thereby managing the interaction behavior rather than following predefined trajectories. This is especially relevant to guarantee safety of the human operator when interacting with a robot.

Compliance can be passive–or mechanical–when the mechanical properties of an actuator or another robotic part are tuned to determine what stiffness or damping it can adopt, thereby adapting to the force applied by a human. The introduction of soft and elastic materials on robot bodies, which prevents harm to humans in case of impact, can be seen as an example of better passive compliance. Active–or cognitive–compliance, on the other hand, implies the use of algorithms to actively model how the stiffness or the damping must change as an effect of task requirements[6].

Advancements in sensors, both for forces and torques, as well as the availability of tactile signals collected by artificial skin[7] played a crucial role in achieving cognitive compliance of robotic systems and making physical human-robot interaction safer. As a result, compliant control has allowed humans to control the robot movements with touch, for instance through haptic interfaces. It has increased precision and performance in tasks execution, as well as safety. Among the compliance control strategies, of particular importance was the development of variable impedance actuators[8], which was based on the intuition of bringing intelligence into the robot hardware.

Recent research in robotics has also addressed social compliance, that relates to the capacity of a system to conform to social norms and expectations. In particular, many works have studied how monitoring physiological signals related to social compliance could allow robots to change their actions, for instance stopping a task and waiting until those signals indicate that the human has regained some of the required trust before resuming the joint work[9].

**Non-Physical HRI**. In non-physical human-robot interaction, better computer vision algorithms have contributed significantly to improved navigation in various environments, to detect humans, to predict their

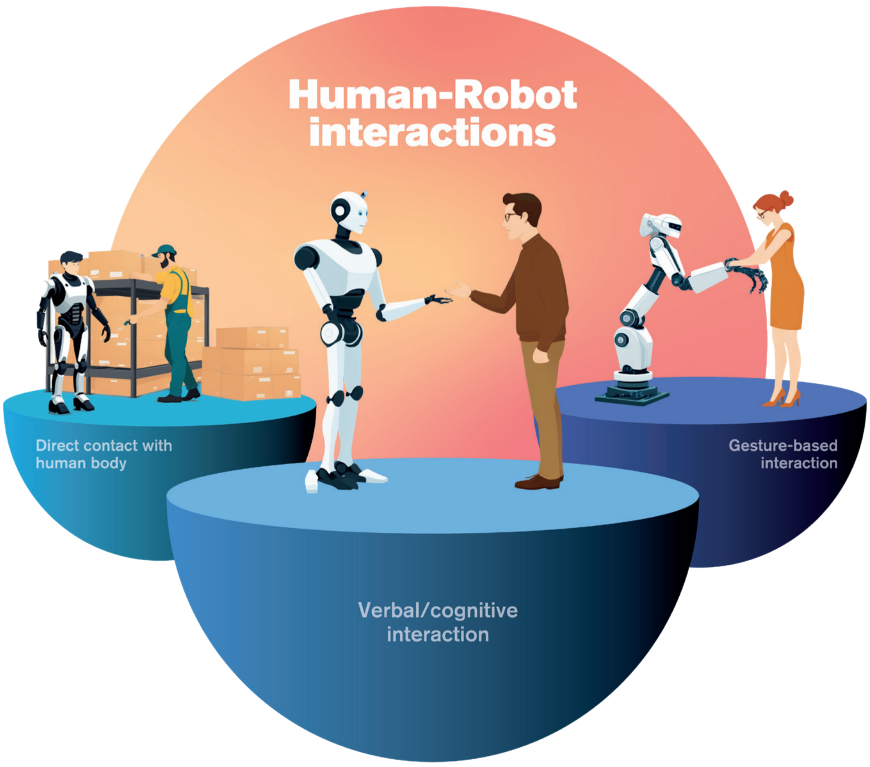


*Fig. 1. Different types of human-robot interaction*

intentions and behave appropriately. Simultaneous localization and mapping (SLAM) is now a mature technology deployed in a wide variety of mobile robots, such as drones and vacuum cleaners. It can detect human as well as non-human obstacles. Computer vision has also helped robots to improve their capabilities in manipulation tasks that involve interactions with humans[10].

Breakthroughs in speech recognition, thanks to recurrent neural networks and transformers, fueled the diffusion of voice assistants, such as Amazon's Alexa, Google Assistant and Apple's Siri. The latest developments in LLMs have boosted verbal human-robot interaction, allowing robots to conduct complex dialogues with humans. However, there is still a stark contrast between the advanced conversational abilities of LLMs and the still limited capacity of robots to interact with the physical world. A language model may allow a robot to engage with humans in a complex conversation on how to set up a dinner table for friends versus hosting a boss, but if the human were to ask the robot to fetch a glass from the table, the robot may not be able to identify the correct glass, or may prove much more clumsier and may knock over other objects along the way.

## 3. COMMERCIAL APPLICATIONS OF HRI

All together, these three decades of research have led to the development of robots designed and programmed to be intrinsically safe for humans—that is, capable of working safely near them without being cognitively aware of their presence. Several applications have emerged from this achievement.

Cobots started operating in the manufacturing industry outside of confined spaces[11]. The first industrial cobot to reach the market was the LBR iiwa single arm system, developed by the German Aerospace Center (DLR) and commercialized by KUKA in 2008[12]. It enabled force and torque sensing in all joints. Other single- or dual- arms systems became available in the following years, such as the UR5 by Universal Robots, Robonaut by NASA, that started operating on the International Space Station, and Baxter by Rethink Robotics, later evolved in a single-arm version called Sawyer. Even if both Baxter and Sawyer had an affordable pricing, they had mixed success. This was maybe due to the adoption of spring actuators coupled with force sensors, which make them safer than cobots that relied on more traditional positioning systems, like those designed by Universal Robots, but less precise[13]. The cobot industry in the end got dominated by Universal Robots single-arm systems. Competitors are emerging in Europe, such as Europe Technologies and Agile Robots.

Robots that navigate and share space safely with humans have been deployed satisfactorily in many environments. Autonomous vehicles have been deployed in factories and warehouses, as in the case of Carter developed by Robust.AI[14], but their interaction with human workers is still carefully structured. Domestic robots performing household tasks, such as vacuum cleaners, are produced by the million each year. Mobile robots work quite well in public spaces such as hospitals and airports. Robots have also been deployed in restaurants to serve people, where they can navigate without ground signaling, finding their way among customers and servers.

A significant leap forward in terms of robots that physically interact with humans has been the development of wearable robotics, especially exoskeletons. Thanks to the development of lighter and more robust materials, such as titan, exoskeletons are now commercially deployed, not only to help physically impaired people but also to assist humans in heavy duties, reducing the social cost of work[15].

One step beyond there are those robots that, thanks to the very anthropomorphic or biomorphic design of their body and of the controller, can interact with humans and be human-aware cognitively. For example, they can talk and respond to humans, simulate facial expressions and recognize human's expression, predict human motion and adapt to it. Robots for rehabilitation and companionship have been deployed and their effectiveness has been demonstrated to a certain extent. One of the most successful examples is PARO, the baby seal robot, whose deployment has been proved useful in hospitals and elderly homes, especially with people affected by dementia[16]. Other examples in this category are CLEO and AIBO, a robot cat and dog respectively which are used mainly with children with well-documented benefits. ROBOTA, an imitating doll robot and Keepon, a small friendly ball-like robot, were both successfully used to engage with children in autism research. Social robots have been developed also with highly realistic human faces, such as Geminoid by Hiroshi Ishiguro Laboratories and Sophia by Hanson robotics. Robots capable of interacting with humans have also been developed for educational purposes. An example is the Sphero robot, employed in programming classes for children.

As for the more theoretical work on human-robot interaction, that is the study of how humans interact with robots and of certain aspects of human behavior which employ robotic platforms, the most robust findings concern how people interact with robots in controlled laboratory settings, one-to-one interaction with humans trained to interact with the robot. One of the pioneering works in this field of research is Kismet[17], the robotic platform developed at MIT in the late 90s. Kismet is a social robot designed to engage in natural and expressive face-to-face interaction with a human. It was inspired by infant social development, behavior and psychology. The human and the robot are led to interact like in a parent-infant relationship. Around 2010, several new robotic platforms for social human-robot interaction emerged, such as NAO, iCub, Kaspar and Pepper. NAO has been among the most widely used social robots in human-robot interaction research due to its affordability and broad functionality. It has been used in various applications, such as education, autism therapy and elderly care[18].

In the last three decades, research on human-robot interaction has achieved important results, bringing robots to interact with humans in different contexts. However, the goal of developing fully autonomous robots capable of interacting usefully and pertinently with humans is still quite far away. Next, we offer our view on the most pressing challenges which HRI must resolve, and the most rapid new deployments of HRI we can expect.

## 4. SHORT-TERM CHALLENGES AND LOW-HANGING FRUITS

HRI is crucial to scaling up the use and deployment of collaborative robots, that is robots capable of operating outside the confined environments of industrial settings—where they traditionally worked in cages or behind fences to prevent any interaction with hu-


*Fig 2. Robots interacting and collaborating with humans. A: Collaborative human-robot sawing experiment performed at Italian Institute of Technology in 2016 (Credit: Luka Peternel, CC BY-SA 4.0). B: A journalist speaking to humanoid robot Sophia at the Deutsche Welle Global Media Forum in 2019 (Credit: Deutsche Welle, CC BY-NC 2.0). C: Pepper robot interacting with a waiter in a Tokyo cafe, 2019 (Credit: International Labour Organization/K. Hongladarom, CC BY-NC-ND 2.0). D: PARO therapeutic seal robot interacting with an elderly woman in a nursing home in 2012 (Credit Amber Case, CC BY-NC 2.0). E: AIBO ERS-7 following pink ball held by child (Credit: Stuart Caie, CC BY 2.0). F: Child interacting with Robota during a behavioral study conducted in 2007 (reproduced with permission from[19]).*

mans. It is also needed to expand usage of robots in the medical sector, wearable robots and exoskeletons, robots and drones for the inspection of remote locations and for search and rescue operations in collaboration with humans, mobile robots capable of navigating in crowded spaces, such as hospitals, airports, restaurants, and robots for social companionship. Ultimately, the most challenging application is inside homes, which are among the most unstructured and unpredictable environments.

Even if the market for collaborative robots in the industry has been growing dramatically in the last ten years, they still represent only 5 to 8% of the robots sold to industries. Their presence in manufacturing and logistics could increase soon as they become safer to interact with, more robust and capable of performing highly dynamic motions similar to what humans do[20].

The physical interaction with robots could be enhanced by providing the robot with tactile sensors that can measure more accurately the contact forces[21]. The idea of an artificial skin has been proposed a long time ago, but so far it has been implemented mainly in robotic platforms for research. One of the first examples was the iCub robot developed at the Italian Institute of Technology that already ten years ago was covered with 200 tactile sensors, one the largest implementations of tactile skin at the time. More recently, researchers at TU-Munich implemented[22] tactile skin based on off-the-shelf components integrated on hexagon shaped printed circuit boards, with which they covered a full-size humanoid robot (H1). One of the key elements was to reduce the computation in a way that allowed the humanoid robot to operate autonomously without needing additional computation or external energy source. Deploying robots with tactile

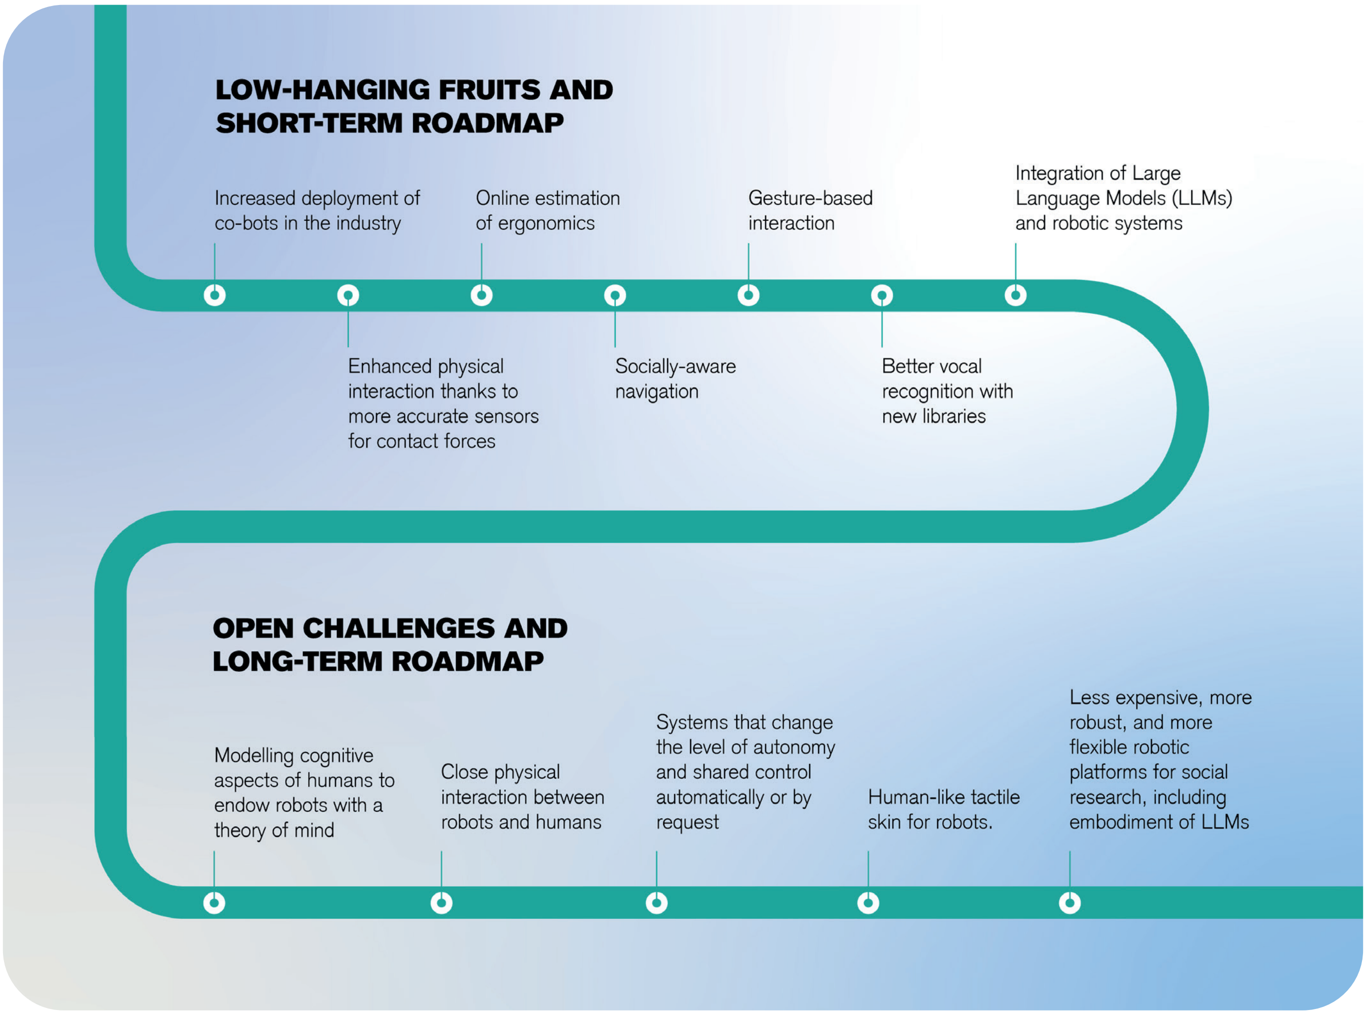


*Figure 1: Short-term and long-term research goals on human-robot interaction*

sensing small covering surfaces is achievable in the short term and could be driven also by the need to improve the manipulation performance of robots employed in the agrifood and textile industries.

Research on online estimation of ergonomics has been extremely active and industrial applications could be achieved soon, facilitated by the fact that the technology works independently of the context. Recently there have been several attempts to offer solutions based on wearable inertial sensors and pressure or force sensors rather than systems based on optical motion capture, which can be easily occluded in cluttered environments such as factories[23 24]. Also, researchers have been working on improving the intuitiveness of ergonomic evaluation and visualization tools based on digital human models to facilitate their adoption by industrial operators[25]. These technologies could be employed in devices that alert workers performing heavy tasks about damaging postures to prevent musculoskeletal disorders and enhance safety in the workplace.

Significant progress on socially-aware navigation is now within reach for the research community, especially in Europe[26]. In the next few years, there could be robots capable of navigating among humans, not only maintaining appropriate distance from humans, but also understanding how humans move when they are confused. Deploying this kind of robots will require refinement of the algorithms that allow the robots to detect and perceive humans, and also predict what will happen in the next second. This would help find a trade-off between safety and usefulness of robots navigating among humans. Socially-aware navigation would also greatly benefit from a better understanding of the interaction between naive users and robots in unstructured settings. As an example, last-mile delivery robots need to share the sidewalk with humans, and more work is needed to model humans' expectations regarding the behavior of these robots[27].

Gesture-based interaction is quite mature now, as there are many algorithms that work well in laboratory conditions. More research is needed to make them work in the real-world environment, where humans do not act perfectly and there could be occlusions and disturbances, but the science is now solid.

Vocal and audio recognition is mature too, with several well-performing libraries. Deploying it could be slightly more challenging than gesture recognition because of the great variance with which people speak, and of course each application can have a different vocabulary. To address real-world acoustic conditions, the signal processing community offers numerous libraries capable of denoising audio signals. For instance, these libraries can effectively extract speech from a drone's onboard microphone despite high levels of ambient noise.

The integration of LLMs and robotic systems holds promise to transform the human-robot interaction paradigm, allowing robots to act upon high-level instructions expressed in natural language and generate plans in the form of step-by-step procedures or code. This field of research is just starting but it is rapidly evolving, with several examples of this integration already available[28]. These include the models developed by Google DeepMind[29] or LaMi by the Honda Research Institute Europe[30]. LaMi converts various forms of human input, such as behavior, position, gaze, dialogue, and scene information, into a language that the LLM can process. The LLM then analyzes the situation and determines how and when the robot should assist humans, following predefined guidelines. Additionally, it synchronizes the robot's movements (lid, neck, ears) with speech output to create dynamic, multi-modal expressions.

As for the more theoretical work on human-robot interaction, researchers will need to understand much more about the interaction of robots with multiple users. This last setting can be used to study how discrimination and social exclusion[31] arise and how biases in various robot components can affect the group members depending on their features[32]. Examples of these so-called "embodied biases" are natural language understanding models and voice recognition models that are known to be better at recognizing male voices than female ones, or image recognition models which are better with white people than with people of color. This is largely due to the underrepresentation of females and people of color in the datasets used to train these models. All these elements will be embedded in robotic platforms and biases can arise in multiple ways[33].

There is also research about the possibility to mitigate some of these risks, making robots aware of discrimination but also that other humans can discriminate. Besides making robots aware of this risk, one could program them to apologize when they are discriminating or if they are at risk of discriminating, explaining why they are doing so to start some process to reintegrate a person in the group.

## 5. OPEN SCIENTIFIC CHALLENGES

The long-term goal of human-robot interaction research is to design robot systems and AI systems which explicitly consider the human in terms of their actions, their preferences, their mental state and their goals and therefore understand when they need to act or communicate. All these questions are far from being solved today.

One open challenge is modeling basic cognitive aspects of humans, to endow robots with a theory of mind that allows them to understand the human's expectations during a joint task and to engage in a negotiation which leads to results that align with the human's preferences, objectives and values[34].

Robots should also be able to update these models with time. As an example, they should understand when a human is not feeling well and performing with less dexterity and pitch in to help, but they should take themselves back as soon as the human recovers. They should be able to do this in open environments where new people may come into play and take roles in the joint actions.

To achieve human-centered robotics, researchers should strive to develop robots that are easy to interact and work with and do not overly constrain humans. In logistics, where robots are already deployed, this kind of problem has already arisen, leading to a high degree of turnover among human workers. These workers often feel replaceable as they work to complement robots. The development of future cobots should be centered around human workers to ensure they do not refrain from intervening with robots out of fear of the consequences this could have on their jobs. The robots of the future should be able to adapt and give priority to the human, allow them the freedom to make their own decisions, and assist rather than impose their rhythm. It is thus crucial to involve experts from fields such as psychology, economics, philosophy, and cognitive science to understand thoroughly what it means to collaborate with humans[35]. European institutions should weigh in with meaningful regulations to enforce the principle of human-centered robotics, as they have already done concerning the use and exploitation of personal data and the deployment of AI systems. Moreover, robots should be well integrated with existing infrastructures, also digital ones.

To fulfill this long-term vision several more specific challenges should be addressed.

In physical human-robot interaction, a challenge will be that of designing systems that allow to change the level of autonomy and shared control. This could happen automatically or by request, so that it fits exactly what the human wants and needs in terms of ergonomics, but especially what the human feels comfortable with. This would allow the human to preserve its own sense of agency and not feel completely dominated by the robot. Solving this challenge will be crucial for exoskeletons, especially active ones, which will be more and more deployed in assisting humans in various tasks, such as lifting heavy payloads, and for manipulators on a mobile base in a factory.

Tactile human-robot interaction would be enhanced by developing human-like tactile skin for robots. One of the necessary steps in this direction is to build electronics and algorithms capable of locally processing the vast amount of sensing data collected over large surfaces. This would reduce the quantity of data processed by the central processing unit, which should only handle high-level decisions about perception.

Currently, AI and machine learning algorithms in robotic sensing are implemented using digital electronics. A new computing paradigm inspired by the human brain, which is based on analog signals, is required. This could be achieved by developing neuromorphic devices—hardware that implements neuromorphic arrangements and is capable of learning, similar to how synapses between neurons create plasticity in the brain. An example is printed synaptic transistors placed close to the sensor that can learn[36].

Other problems that need to be solved to have humanoid robots performing a wide variety of tasks in collaboration with humans concern the ability to learn in a continuous way, and the ability to personalize behavior to different people. Also, such robots need to be soft enough. A major development step is required also on all levels of hardware and control to enable close physical interaction, for tasks like feeding or undressing, washing and dressing again a person, currently performed by caregivers. Humanoid robots can play multiple roles because they fit in the environments we have built for us humans, both in terms of shape and size. Moreover, getting assistance from a humanoid robot is more effective because humans know how to behave in such situations, e.g. they know they can put a hand on the shoulder or how to walk together.

As the interaction between robots and humans becomes more unconstrained, more interdisciplinary research is needed, especially on managing humans' expectations of the robots' capabilities.

To investigate the social aspects of human-robot interaction, cheaper, more robust, and more flexible robotic platforms are needed. Currently, the choice is between robot toys, which are very robust and very cheap but with a very specific application area which cannot be changed easily, and research platforms, such as Pepper and NAO, which are instead quite expensive. In between there are platforms developed by computer scientists and engineers that are often too complex to customize. As a result, the choice of the platform is rather restricted. Today, a long-term study on social robots requires major investments. Real-world data on long term interaction between robots and humans is not yet available, since companies like Jibo or Blue Frog Robotics that aimed to develop robot companions did not take off[37]. Their products didn't make the step into people's homes as expected.

AI and machine learning techniques, especially deep learning and LLMs will play an important role in developing the robots of the future. The integration of LLMs into robotic systems could enable individuals, irrespective of their technical knowledge, to interact with robots and direct their actions. An essential step forward compared to LLMs today will be the capability to generate safe and reliable actions in the physical world, based on a physical architecture which is fully aware of the robot's internal state and capabilities. This would require at the same time to build versatile robots, which are capable of doing many different things, navigate, pick objects, interact safely with the environment, avoid obstacles, etc. It is a huge endeavor, but it will change the way in which humans interact with robots.

The use of LLMs could prove highly effective in social science research. For example, studies on the impact of language tone in communication, both between humans and between robots and humans, are challenging to perform by asking humans to change their tone, whereas LLMs can do this very efficiently. However, attention must be paid to reproducibility, especially in long-term studies. Given that LLMs evolve rapidly due to the availability of new training data, their consistency over the course of experiments or across experiments performed at different times should not be taken for granted.

However, as for other fields of robotics, the challenge in human-robot interaction lies in finding an intelligent way to combine model-based AI systems with deep learning algorithms, to mitigate potential risks such as misinterpretation. This requires defining in which situations misinterpretation can be accepted, because it poses no safety issues, and situations where we need instead that the machine really understands what happened, to assess it correctly. As of today, we cannot rely on LLMs for this, and we need to complement them with conservative measures to avoid dangerous consequences.

Embodiment of LLMs will also require dealing with safety and trust, since people are going to be able to use and interact with the robots even without knowing how they work and what their limits are. The integration of ChatGPT into a robotic arm that collaborates with human workers on an assembly task has demonstrated that it significantly increased trust in human-robot collaboration[38]. This is an opportunity, but it also poses the risk of over-trusting, as it has been observed in autonomous cars where, even with the vehicle alerting them that its sensors are malfunctioning and asking them to take over, people did not intervene, trusting that the car will recover. Strategies to recognize excessive trust and refusing to execute dangerous plans in such situations should be developed and deployed.

# 6. CLOSING WORDS

To fulfill the vision of robots interacting with humans in unstructured environments and collaborating with them to perform a wide variety of tasks, human-robot interaction research is becoming increasingly central to robotics. The first steps toward that vision were made over the last three decades, thanks to compliant control, both mechanical and cognitive, new sensors for vision and touch, and progress in voice recognition and natural language processing, thay has enhanced the complexity of dialogues between humans and robots. As a result, cobots entered factories and started navigating public spaces, and several platforms are now employed in the treatment of neurodevelopmental disorders and neurodegenerative diseases, as well as for companionship. However, as of today interaction is still quite constrained and where it is more widespread, such as in logistics, humans often perceive robots as imposing their rhythm rather than adapting to human needs. To progress further towards human-centred robotics, it is crucial to conduct research on managing autonomy levels and understanding humans' expectations and preferences. Additionally, advancing tactile sensing will be critical if robots are to help humanity to tackle the societal challenges it is facing, such as supporting an aging population and mitigating the impacts of climate change. Finally, the integration of AI and machine learning into robotics promises to make robots more accessible to people without technical expertise. While this opens up new perspectives, it also entails risks that need to be addressed. Humans might over-trust robots, underestimating potential hazards, or fall victim to embodied biases—discriminatory behaviors stemming from imbalanced training data used in AI systems.

# REFERENCES


[1] V. Vatsal and G. Hoffman, “Wearing your arm on your sleeve: Studying usage contexts for a wearable robotic forearm,” 2017 26th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN), Lisbon, Portugal, 2017, pp. 974-980, doi: 10.1109/RO-MAN.2017.8172421

[2] M. L. Elwin, B. Strong, R. A. Freeman and K. M. Lynch, “Human-Multi-robot Collaborative Mobile Manipulation: The Omnid Mocobots,” in IEEE Robotics and Automation Letters, vol. 8, no. 1, pp. 376-383, 2023, doi: 10.1109/LRA.2022.3226366.

[3] H. Admoni and B. Scassellati, “Social eye gaze in human-robot interaction: a review,” J. Hum.-Robot Interact., vol. 6, no. 1, pp. 25–63, 2017, doi: 10.5898/JHRI.6.1.Admoni.

[4] P. D. H. Nguyen, F. Bottarel, U. Pattacini, M. Hoffmann, L. Natale and G. Metta, “Merging Physical and Social Interaction for Effective Human-Robot Collaboration,” 2018 IEEE-RAS 18th International Conference on Humanoid Robots (Humanoids), Beijing, China, 2018, pp. 1-9, doi: 10.1109/HUMANOIDS.2018.8625030.

[5] M. Schumacher, J. Wojtusch, P. Beckerle and O. von Stryk, “An introductory review of active compliant control,” Robotics and Autonomous Systems, vol. 119, pp. 185-200, 2019. doi: https://doi.org/10.1016/j.robot.2019.06.009.

[6] A. Billard, Robotics and Autonomous Systems 88, 157-164 (2017). DOI: 10.1016/j.robot.2016.08.030

[7] G. Cheng et al. Proceedings of the IEEE 107(10), 2034-2051(2019). DOI: 10.1109/JPROC.2019.2933348

[8] B. Vanderborght et al., “Variable impedance actuators: A review,” Robotics and Autonomous Systems, vol. 61, no. 12, pp. 1601-1614, 2013, doi: 10.1016/j.robot.2013.06.009.

[9] A. Billard, “ On the mechanical, cognitive and sociable facets of human compliance and their robotic counterparts,” Robotics and Autonomous Systems vol. 88, pp. 157-164, 2017, doi: 10.1016/j.robot.2016.08.030.

[10] F. Cini V. Ortenzi, P. Corke and M. Controzzi, “On the choice of grasp type and location when handing over an object,” Science Robotics, vol. 4, p. eaau9757, 2019, doi: 10.1126/scirobotics.aau9757.

[11] A. Hentout, M. Aouache, A. Maoudj and I. Akli, “ Human–robot interaction in industrial collaborative robotics: a literature review of the decade 2008–2017,” Advanced Robotics, vol. 33, no. 15–16, pp. 764–799, 2019, doi: 10.1080/01691864.2019.1636714.

[12] R. Bischoff et al., “The KUKA-DLR Lightweight Robot arm - a new reference platform for robotics research and manufacturing,” ISR 2010 (41st International Symposium on Robotics) and ROBOTIK 2010 (6th German Conference on Robotics), Munich, Germany, pp. 1-8 (2010). Accessed: Dec. 10, 2024. [Online]. Available: https://ieeexplore.ieee.org/document/5756872

[13] C. Lawrence, “Rise and Fall of Rethink Robotics,” The American Society of Mechanical Engineers, 2019. Accessed: Dec. 10, 2024. [Online]. Available: https://www.asme.org/topics-resources/content/rise-fall-of-rethink-robotics

[14] Robust AI Blog. Accessed: Dec. 10, 2024. [Online]. Available: https://www.robust.ai/blog

[15] M. Bär, B. Steinhilber, M.A. Rieger and T. Luger, “The influence of using exoskeletons during occupational tasks on acute physical stress and strain compared to no exoskeleton – A systematic review and meta-analysis,” Applied Ergonomics, vol. 94, p. 103385, 2021, doi: 10.1016/j.apergo.2021.103385.

[16] H.S. Kang, K. Makimoto, R. Konno, I.S. Koh, “Review of outcome measures in PARO robot intervention studies for dementia care,” Geriatric Nursing, vol. 41, no. 3, pp. 207-214, 2020, doi: 10.1016/j.gerinurse.2019.09.003.

[17] C. Breazeal, “Sociable Machines: Expressive Social Exchange Between Humans and Robots”. Sc.D. dissertation, Department of Electrical Engineering and Computer Science, MIT (2000). Accessed: Dec. 10, 2024. [Online]. Available: https://groups.csail.mit.edu/lbr/hrg/2000/phd.pdf

[18] A. Amirova, N. Rakhymbayeva, E. Yadollahi, A. Sandygulova and W. Johal, “10 Years of Human-NAO Interaction Research: A Scoping Review,” Frontiers in Robotics and AI, vol. 8, 2021, doi: 10.3389/frobt.2021.744526.

[19] A. Billard, B. Robins, J. Nadel and K. Dautenhahn, “Building Robota, a Mini-Humanoid Robot for the Rehabilitation of Children With Autism,” Assistive Technology, vol. 19, no. 1, pp. 37–49, doi: 10.1080/10400435.2007.10131864.

[20] M. Grebenstein et al., “The DLR hand arm system,” 2011 IEEE International Conference on Robotics and Automation, Shanghai, China, 2011, pp. 3175-3182, doi: 10.1109/ICRA.2011.5980371.

[21] B. D. Argall and A. G. Billard, “A survey of Tactile Human–Robot Interactions,” Robotics and Autonomous System, vol. 58, no. 10, pp. 1159-1176, 2010, doi: 10.1016/j.robot.2010.07.002.

[22] S. Groß et al., “Soft Sensing Skins for Arbitrary Objects: An Automatic Framework,” 2023 IEEE International Conference on Robotics and Automation (ICRA), London, United Kingdom, 2023, pp. 12507-12513, doi: 10.1109/ICRA48891.2023.10161344.

[23] P. Maurice et al., “Human movement and ergonomics: An industry-oriented dataset for collaborative robotics,” The International Journal of Robotics Research, vol. 38, no. 14, pp. 1529-1537, 2019, doi: 10.1177/0278364919882089.

[24] S. Helmstetter and S. Matthiesen, “Human Posture Estimation: A Systematic Review on Force-Based Methods—Analyzing the Differences in Required Expertise and Result Benefits for Their Utilization,” Sensors, vol. 23, no. 21, p. 8997, 2023, doi: 10.3390/s23218997.

[25] L. Vianello, W. Gomes, F. Stulp, A. Aubry, P. Maurice and S. Ivaldi, “Latent Ergonomics Maps: Real-Time Visualization of Estimated Ergonomics of Human Movements,” Sensors, vol. 22, no. 11, p. 3981, 2022, doi: 10.3390/s22113981.

[26] P. T. Singamaneni et al., “A Survey on Socially Aware Robot Navigation: Taxonomy and Future Challenges,” The International Journal of Robotics Research, vol.43, no.10, 2024, doi:10.1177/02783649241230562.

[27] A. M. H. Abrams et al., “Field observation: Interactions between pedestrians and a delivery robot”, 29th IEEE International Conference on Robot & Human Interactive Communication. Accessed: Dec. 10, 2024.

[Online]. Available: https://crowdbot.eu/wp-content/uploads/2020/09/Short-Talk-1-Workshop_Abstract_Field-Observation_final.pdf

[28] C. Zhang, J. Chen, J. Li b, Y. Peng and Z. Mao, “Large language models for human–robot interaction: A review,” Biomimetic Intelligence and Robotics, vol. 3, no. 4, p. 100131, 2023, doi: 10.1016/j.birob.2023.100131 .

[29] Google DeepMind, “Demonstrating Large Language Models on Robots”, Robotics: Science and Systems 2023. Accessed: Dec. 10, 2024. [Online]. Available: https://www.roboticsproceedings.org/rss19/p024.pdf

[30] C. Wang et al., “LaMI: Large Language Models for Multi-Modal Human-Robot Interaction,” 2024 CHI Conference on Human Factors in Computing Systems, Article 218, 1-10, doi: 10.1145/3613905.3651029.

[31] S. Mongile et al., “What if a Social Robot Excluded You?: Using a Conversational Game to Study Social Exclusion in Teen-robot Mixed Groups,” Proceedings of HRI’23 Companion, pp. 208-212, 2023, doi: 10.1145/3568294.3580073.

[32] A. Rosenthal-von der Pütten and A.M.H. Abrams, “Social Dynamics in Human-Robot Groups – Possible Consequences of Unequal Adaptation to Group Members Through Machine Learning in Human-Robot Groups,” . In: Degen, H., Reinerman-Jones, L. (eds) Artificial Intelligence in HCI. HCII 2020. Lecture Notes in Computer Science, vol 12217. Springer, Cham. DOI: 10.1007/978-3-030-50334-5_27.

[33] A. Rosenthal-von der Pütten and S. Schiffer, “Potential Ways to Detect Unfairness in HRI and to Re-establish Positive Group Dynamics,” arXiv [cs.RO], 2023, doi: 10.48550/arXiv.2310.01574.

[34] S. Lemaignan, M. Warnier, E. Akin Sisbot, A. Clodic and R. Alami, “Artificial cognition for social human–robot interaction: An implementation,” Artificial Intelligence, vol. 247, pp. 45-69, 2017, doi: 10.1016/j.artint.2016.07.002.

[35] D. Kragic and Y. Sandamirskaya, “Effective and natural human-robot interaction requires multidisciplinary research,” Science Robotics, vol. 6, no. 58, p. eabl7022, 2021, doi: 10.1126/scirobotics.abl7022

[36] F. Liu et al., “Printed synaptic transistor–based electronic skin for robots to feel and learn,” Science Robotics, vol. 7, no. 67, p. eabl7286, 2022, doi: 10.1126/scirobotics.abl7286.

[37] J. van Camp, “Review: Jibo Social Robot. This robot roommate holds a lot of promise but still has a long way to go.”, Wired, 2017. Accessed: Dec. 10, 2024. [Online]. Available: https://www.wired.com/2017/11/review-jibo-social-robot/

[38] Y. Ye, H. You and J. Du, “Improved Trust in Human-Robot Collaboration with ChatGPT,” arXiv [cs.RO], 2023, doi: 10.48550/arXiv.2304.12529.